



## **Infoblox NIOS 8.6.x Documentation**

07/11/2023

# Contents

<b>What's New .....</b>	<b>3</b>
<b>Support Matrix .....</b>	<b>15</b>
<b>Getting Started .....</b>	<b>16</b>
<b>Installing NIOS .....</b>	<b>30</b>
<b>Upgrading NIOS .....</b>	<b>32</b>
<b>Using the Grid Manager Interface .....</b>	<b>64</b>
<b>Administering NIOS .....</b>	<b>183</b>
<b>Using the NIOS CLI .....</b>	<b>1940</b>
<b>Using NIOS APIs .....</b>	<b>2167</b>
<b>Reference Information .....</b>	<b>2168</b>

## What's New

The NIOS 8.6 release includes the following new features and enhancements:

### Bypassing Subscriber Secure Policy for Allowed Lists (RFE-11652)

NIOS 8.6.3 introduces the **Enable Subscriber Secure Policy Bypass for Allowed list** and the **Set Global Allow List RPZ index range ( 0 to 30)** options in which subscriber specific allow domains take priority over category based policies (content-filtering), security policies and blocklist entries. Subscriber specific block domains take priority only over category-based policies (content-filtering).

The **Enable Subscriber Secure Policy Bypass for Allowed list** checkbox enables NIOS to generate a normal response for all domains in a subscriber's allowed list. Allowed domains will override RPZ rules if any (for example, NXDOMAIN), and categorization policy rules for the subscriber. This enables subscribers to override all policies for a specific domain. The **Set Global Allow List RPZ index range ( 0 to 30)** checkbox adds a domain to an RPZ specified as a passthru RPZ rule, and that domain is also added as a global allowed list. NIOS 8.6.3 also introduces a new report called *Daily Report on Subscriber Information and DNS Queries matching with ABLs* that displays the overall subscriber information and DNS domain queries that match the queried domains in the allow and block lists. For more information about these options, see the [Scaling Subscriber Sites](#) topic in the NIOS 8.6 online documentation.

### Validating Certificates (SPTYRFE-52)

NIOS 8.6.3 introduces the **Grid > Certificates > Validate Certificates** option that validates all certificate uploaded using the Manage Certificates option and also validates DNS Traffic Control HTTPS health monitor certificates. It displays the certificates to be either valid, invalid, or expired. It displays warning messages for expired certificates and for certificates with no SKI (Subject Key Identifier). For more information, see the [Validating Certificates](#) topic in the NIOS 8.6 online documentation.

### ZVELO Category Database Update Failure Changes (RFE-12140)

If a ZVELO category database update failure occurs for three consecutive days:

- Grid Manager displays a yellow background with the "Please correct the download credentials or the proxy configuration to get the latest database updates" message and the member status is displayed as "Domain category db is not latest" in the **Grid Manager > Subscriber Collection > Services > Service Status** column.
- A new SNMP trap is sent with the message "Domain category db is not latest". Additionally, if email notifications are configured, an email is sent to the configured email address with the "Domain category db is not latest" message.
- Post this event, if the ZVELO download is successful, a new SNMP clear trap is sent, and an email with the "zvelo SNMP Clear Trap" message is also sent. The Service Status column has a green background and the "Subscriber service is working" message is displayed.

If a ZVELO category database update failure occurs for more than 60 days:

- Grid Manager displays a red background with the "zvelo database expired. Subscriber secure queries will be fail-open." message and the member status is displayed as "zvelo db has expired" in the **Grid Manager > Subscriber Collection > Services > Service Status** column.
- A new SNMP trap is sent with the "zvelo db has expired" message. Additionally, if email notifications are configured, an email is sent to the configured email address, with the "zvelo db has expired" message.
- Post this event, if the zvelo download is successful, a new SNMP clear trap is sent, and an email with the "zvelo SNMP Clear Trap" message is also sent. The Service Status column has green background and the "Subscriber service is working" message is displayed.

## Support for Network Insight on the AWS and Microsoft Azure Public Cloud (RFE-10248)

Starting from NIOS 8.6.3, deployment of Network Insight appliances on Microsoft Azure and AWS public cloud is supported. You can use the discovery feature to detect devices in your network and manage device data from Grid Manager.

## Filtering Options for vDiscovery

Starting from NIOS 8.6.3, you can configure vDiscovery jobs with CIDR-based filters to limit vDiscovery only to configured networks or to skip vDiscovery from excluded networks. For more information, see the [Configuring vDiscovery Jobs](#) topic in the the NIOS 8.6 online documentation.

## Infoblox BloxConnect Changes

From NIOS 8.6.3 onwards, the Infoblox BloxConnect screen that appears when you first log in to Grid Manager will no longer be displayed. To better enhance the NIOS customer experience and for Infoblox to proactively respond to customer issues in a timely manner, BloxConnect data will be collected by default and the same data will be included in the support bundle. For information about the data collected, see <https://insights.infoblox.com/resources/datasheets/infoblox-datasheet-infoblox-bloxconnect>.

## Synchronizing Amazon Route 53 Data from Multiple AWS Accounts (RFE-9382)

vNIOS for AWS is integrated with Amazon Route 53. You can now discover and synchronize Amazon Route 53 DNS data from multiple AWS accounts of an AWS organization to NIOS using a single NIOS admin account. For more information, see the Amazon Route 53 topic in the *Installation Guide for vNIOS for AWS* at <https://docs.infoblox.com>.

## Support for Route 53 Integration on AWS GOV Cloud (RFE-11806)

Starting from NIOS 8.6.3, the vNIOS for AWS Route 53 integration is supported on AWS GOVCloud. For more information, see the Amazon Route 53 topic in the *Installation Guide for vNIOS for AWS* at <https://docs.infoblox.com>.

## Support for Amazon EC2 R6i Instance Type (RFE-12312)

vNIOS for AWS instances running on NIOS 8.6.3 or later can be deployed on Amazon EC2 R6i instance types. For information about supported EC2 shapes and models, see the Infoblox vNIOS for AWS AMI Shapes and Regions topic in the *Installation Guide for vNIOS for AWS* at <https://docs.infoblox.com>.



## Connecting to the Serial Console of an Amazon EC2 Instance (RFE-11803)

You can now connect to the serial console of a vNIOS for AWS EC2 instance that is deployed on an R6i instance type for troubleshooting purposes. For more information, see the Provisioning vNIOS for AWS Using the BYOL Model topic in the *Installation Guide for vNIOS for AWS* at <https://docs.infoblox.com>.

## Support for EBS Encryption (RFE-11931)

vNIOS for AWS instances running on NIOS 8.6.3 or later versions of 8.6.x, support encryption of Amazon Elastic Block Storage volumes. For more information, see the Provisioning vNIOS for AWS Using the BYOL Model topic in the *Installation Guide for vNIOS for AWS* at <https://docs.infoblox.com>.

## Enabling the SHA1 Encryption Algorithm for the NTP Key (RFE-8178)

In NIOS 8.6.3, you can add the SHA1 NTP authentication key before enabling the NTP service on the Grid. The key is a 40 character hexadecimal string and it uses a hash-based symmetric encryption algorithm. For more information, see the *Using NTP for Time Settings* topic in the NIOS 8.6 online documentation.

## Exporting to CSV Before Deleting a Network (RFE-11846)

NIOS 8.6.3 introduces the **Export & Delete** button that allows you to export the network data into a CSV file before deleting a network. For more information, see the Exporting and Deleting Networks section in the *Configuring IPv4 Networks* topic in the NIOS 8.6 online documentation.

## Reconnecting Groups After a Grid Master Candidate Promotion (RFE-4753)

In NIOS 8.6.3, you can reconnect groups after a Grid Master Candidate promotion thus giving you more control over the promotion and minimizing service outages by allowing you to group Grid members and schedule a time for the groups to reconnect to the newly promoted Grid Master. As soon as the scheduled time arrives, members of Grid Master Candidate groups reconnect to the newly promoted master. You can do this by using the **Activate GMC Group Promotion Schedule** option in the *GMC Group Promotion Schedule* editor. For more information, see the *Managing a Grid* topic in the NIOS 8.6 online documentation.

## Device Handling Performance Optimization for NIOS Subscriber Cache (RFE-12397)

Device handling performance optimization optimizes the use of NIOS subscriber cache by updating only the provisioned devices to NIOS thereby reducing the number of devices delivered to the NIOS cache. For more information, see the **Device Handling Performance Optimization for NIOS Subscriber Cache** section in the *Infoblox Subscriber Insight and Subscriber Policy Enforcement* topic in the NIOS 8.6 online documentation.

## Improvements in Internal Cache Handling (RFE-12003)

Fast replication now does not take place for all subscriber site members; instead, it takes place from a single member with a full cache thus greatly improving cache handling.

## Discovering Juniper Mist Devices

You can now discover Juniper Mist devices using Network Insight. For more information, see the [Configuring Discovery for SDN and SD-WAN](#) topic in the NIOS 8.6 online documentation.

## SHA-512 and SHA-256 Support in Network Insight (RFE-12053)

Network Insight now supports the SHA-512 and SHA-256 algorithms for SNMP polling.

## AES-256 Support for SNMPv3 Authentication (RFE-10304)

From NIOS 8.6.3 onwards, the AES-256 encryption is supported for SNMPv3 authentication in Network Insight.

## Converting Real-Time Alerts to Scheduled Alerts (RFE-12594)

In NIOS 8.6.3, alert types in NIOS Reporting for the new device discovery alert, security alert, and category alert have been converted from *Real-time* to *Scheduled*.

## Enhancements in the DNS Object Count Trend and FLEX Grid Licensing Features Enabled Reports (RFE-11909)

In NIOS 8.6.3, certain changes have been implemented to the following reports:

- Managed DDI features enabled
- SPLA Grid Licensing Features Enabled
- Managed DDI Peak IP Usage Trend
- Managed Trend DNS Peak Usage
- DNS Effective Peak Usage Trend for SPLA Grid License

Changes to these reports include:

- All of these reports now also support non -IB-FLEX appliances.
- All of these reports have the **Members** and the **Reporting SPLA** filters.
- If you want to include a Grid member that does not have an extensible attribute in these reports, you must set the Reporting SPLA extensible attribute to **Managed IB-FLEX** or **Managed HW/SW** depending on the Grid member type, and the Grid license must be **flex\_grid\_ms**.

## Support for Microsoft Server 2022 (RFE-12246)

NIOS 8.6.3 supports Microsoft Server 2022.

## Support for VMware ESXi 7.0.x (RFE-12555)

vNIOS for VMware is now supported on VMware ESX or ESXi versions 7.0.2 and 7.0.3.

## Support for Nutanix AOS 6.x LTS (RFE-12650)

vNIOS for Nutanix™ AHV is now supported on NIOS 8.6.3, NIOS 8.6.2, and NIOS 8.5.5. For more information, see the *Infoblox Installation Guide vNIOS for Nutanix AHV* at <https://docs.infoblox.com>.

## IB-FLEX Support for Nutanix AHV (RFE-12601)

vNIOS for Nutanix™ AHV is now supported on IB-FLEX appliances. For more information, see the *Infoblox Installation Guide vNIOS for Nutanix AHV* at <https://docs.infoblox.com>.

## Virtual Advanced DNS Protection Software and Virtual DNS Cache Acceleration Support in vNIOS for Nutanix (RFE-12615)

vNIOS for Nutanix now supports virtual Advanced DNS Protection Software (vADP).

## DHCP, IPv6, and HA Support in vNIOS for Red Hat OpenShift (RFE-12835)

vNIOS for Red Hat OpenShift can now manage DHCP, IPv6, and HA configurations. For more information, see the *Infoblox Installation Guide vNIOS for Red Hat OpenShift* at <https://docs.infoblox.com>.

## Support for Upgraded Splunk Version 8.2.10

NIOS 8.6.3 supports Splunk version 8.2.10.

## NIOS Enhancements to Support Consolidated Hotfixes

NIOS 8.6.3 introduces new enhancements to the implementation of hotfixes. NIOS 8.6.3 introduces the manifest file that is generated after the hotfix is successfully applied. The manifest file contains key details such as affected files, RPMs installed, and the best suggested action to activate changes. NIOS 8.6.3 also introduces CLI support to view the best suggested action to activate hotfix changes in the Grid member and to support the download of manifest files using WAPI and support bundles.

## Enabling and Disabling DNS Traffic Control Objects (RFE-7088)

NIOS 8.6.2 introduces the *Enable Traffic Management Objects* and the *Disable Traffic Management Objects* screens using which you can enable or disable individual DNS Traffic Control objects. You can also disable the health monitoring of a particular object to stop performing health checks. You can access the new screens from the **Traffic Control** panel on the **Data Management > DNS > Traffic Control** tab. For more information about enabling and disabling DNS Traffic Control objects, see the [Managing DNS Traffic Control Objects](#).

## Multi-Master DNS Failover for DDNS (RFE-5514)

In NIOS 8.6.2, if you have configured more than one Grid DNS primary server for DDNS updates for multi-master zones, DHCP servers use the first available DNS primary server that is configured. If the first DNS primary server is not reachable or is offline, then the DHCP servers reach for the next DNS primary server in the preferred multi-domain DDNS list and so on. You can add up to a maximum of three DNS primary nameservers for each zone.

## Configuring Microsoft Servers and Delegated Name Servers (RFE-10168)

From NIOS 8.6.2 onwards, you will not be able to add a delegated name server group only if DNS synchronization is enabled on any Microsoft server configured in NIOS. You also cannot enable DNS synchronization for Microsoft servers in NIOS if delegated name servers are configured on them.

## Support for Upgraded Splunk Version 8.2.4

NIOS 8.6.2 supports Splunk version 8.2.4.

## vNIOS for Azure and vNIOS for AWS Support in IB-V5005 (RFE-7962)

IB-V5005 support is now extended to vNIOS for Azure and vNIOS for AWS. For detailed information see the *Infoblox Installation Guide for vNIOS for AWS* and the *Installation Guide for vNIOS for Microsoft Azure* available at [docs.infoblox.com](https://docs.infoblox.com).

## vNIOS for GCP Support in IB-V4015 and IB-V4025 (RFE-11349)

IB-V015 and IB-V4025 now support vNIOS for GCP. For detailed information see the *Infoblox Installation Guide for vNIOS for GCP* at [docs.infoblox.com](https://docs.infoblox.com).

## DHCP Support on vNIOS for GCP (RFE-9945)

vNIOS for GCP instances running on NIOS 8.6.2 offer DHCP services for on-premise networks. For more information see the *Infoblox Installation Guide for vNIOS for GCP* at [docs.infoblox.com](https://docs.infoblox.com).

## vNIOS for Nutanix AHV 5.20.3 (RFE-11997)

NIOS 8.6.2 supports the deployment of vNIOS on Nutanix AHV 5.20.3. For more details, see the *Infoblox Installation Guide vNIOS for Nutanix AHV* at [docs.infoblox.com](https://docs.infoblox.com).

## Enabling and Disabling BFD Internal DNS Monitoring (SPTYRFE-49)

NIOS 8.6.2 introduces a new checkbox called **BFD Internal DNS Monitoring** in the *Grid Member Properties* editor > **Anycast** tab. Selecting this checkbox enables the internal DNS monitor to send and receive DNS responses and to retract the OSPF or BGP route if it does not receive a DNS response.

You can enable or disable the **BFD Internal DNS Monitoring** checkbox only if you select the **Enable BFD** checkbox. When you enable the **BFD Internal DNS Monitoring** checkbox, you have the option to toggle between enabling or disabling the internal DNS monitor. When you select this checkbox, Infoblox recommends that you also select the **Enable DNS Health Check** checkbox in the *Grid Properties Editor* or the *Member Properties Editor*. The **BFD Internal DNS Monitoring** checkbox is enabled by default. For more information see the [About BFD](#).

## Support for Cisco ISE Integration Through Outbound Endpoint

NIOS 8.6.2 supports Cisco ISE versions 3.0 and 3.1. Infoblox recommends that you configure Cisco ISE 3.0 and 3.1 using the **Outbound Endpoint** tab. Cisco ISE version 3.1 (pxGrid 2.0) is supported only through the Cisco outbound endpoint. For more information see the [Configuring Outbound Endpoints](#).

## Enforcing the Global Proxy List

In NIOS 8.6.2, if you want to proxy the traffic through the MSP (Multi-Services Proxy) server and have categorized the queried domains in the incoming traffic to the global proxy list, then the query resolves to an MSP virtual IP address and NIOS generates a synthetic resolution. For more information, see the [Scaling Subscriber Sites](#).

## New DNSKEY Algorithm (RFE-6068 and RFE-9845)

You can now add the ECDSAP/SHA-256 and ECDSAP/SHA-384 cryptographic algorithms which the Grid Master can use when it generates the Key-Signing Key Rollover (KSK) and Zone-Signing Key Rollover (ZSK).

## Extensible Attribute-based Topology Rulesets (RFE-9107 and RFE-11133)

You can now specify IPAM objects types, network containers, networks, ranges, and hosts and their External Attribute (EA) values in the **Extensible Attributes Source Types for Topology Rules** field to be used as source types when defining DNS Traffic Control topology rules. For more information, see the [Configuring DNS Traffic Control Properties](#).

## vNIOS Support for Microsoft Azure Stack Hub (RFE-8303)

You can now deploy the NIOS virtual appliance on Microsoft Azure Stack Hub. vNIOS for Microsoft Azure Stack Hub which is a hybrid cloud platform that enables a vNIOS appliance to deliver Azure services in an on-prem environment. You can deploy vNIOS for Azure Stack Hub instances from the Azure CLI or the Azure Stack Hub portal. For more information, see the vNIOS Infoblox Installation Guide for Microsoft Azure at [docs.infoblox.com](https://docs.infoblox.com).

## Health Status of DNS Traffic Control Objects in API Responses (RFE-9893)

The Grid Master Candidate now provides the health status of DNS Traffic Control objects such as servers, pools, and LBDNs through WAPI requests.

## Regenerating the Anycast Password (RFE-11117)

This release of NIOS introduces the `set regenerate_anycast_password` command that regenerates the anycast service password. The regenerated 8-character alphanumeric password is saved to the NIOS database and is used across all anycast configuration files (`ospf.conf/bgp.conf/bfd.conf`) for the following CLI commands: `show ospf`, `show bgp`, `show ipv6_ospf`, `show ipv6_bgp`, `show bfd`.

This command is a maintenance mode command and has no arguments. Only superusers can execute this command. The value of password and enable password in the output of the configuration file commands such as `show bfd` are encrypted when you run the command. For more information, see the [set regenerate\\_anycast\\_password](#).

## Viewing Lightweight Access Point Details in Network Insight (RFE-9556)

You can now view the discovered lightweight access points on the **Data Management > Devices** page. The table displays the following information about the discovered lightweight access points: their name, IP address, device type, model, vendor, and device version. You can also view the discovery statuses and other information in the **Discovery Status** table (**Data Management > Devices > Discovery Status**).

## Displaying the Lead Secondary Column in Name Server Group (RFE-2804)

You can now determine which member is configured as a lead secondary by adding in a column to the **Authoritative Zone > Name Servers** tab.

## Support for creation\_time for Host Records (RFE-8509)

NIOS 8.6.1 introduces the option of adding, updating, listing the creation timestamp value of DNS and non-DNS host records using Grid Manager and WAPI.

## Support for IB-V4015 on Red Hat OpenShift (RFE-11545)

Red Hat OpenShift is now supported on IB-V4015 virtual appliance. For more information, see the Infoblox Installation Guide vNIOS for Red Hat OpenShift at [docs.infoblox.com](https://docs.infoblox.com).

## New Port Placements for the Infoblox 2205 and Infoblox 4005 Series Appliances

The front panels of the Infoblox 2205 Series and the Infoblox 4005 Series have been modified to have slots for the four ports (LAN2, HA, LAN1, MGMT) at the right. However, the Infoblox 2205 and Infoblox 4005 Series models that have the ports located at the center are also being shipped. There is no difference in software functionality between the models that have ports on the right and those that have ports in the center. Both the models will support NIOS versions prior to 8.5.4 and earlier.

For a visual representation of these models, see the Infoblox Installation Guide for 2205 Series Appliances and the Infoblox Installation Guide for 4005 Series Appliances documentation at <https://docs.infoblox.com>

## ACL Support for the Last Queried Time in DNS Scavenging (RFE-7933)

You can now create an ACL or ACE for the **Last Queried Time** field in DNS scavenging and thus prevent a specified set of ACLs or ACEs from updating the last queried timestamp. A new GUI field called **Prevent the following ACLs or ACEs from updating the last queried timestamp** in the **Grid DNS Properties > DNS Scavenging > Basic** tab has been introduced. The set of ACL or ACEs can include IPv4 and IPv6 addresses and networks. For more information, see [DNS Record Scavenging](#).

## New Load Balancer to Add Persistence (RFE-6827)

You now have a new load balancing method called Source IP Hash to configure DNS Traffic Control pools. In this method, requests are distributed based on the hash value of an IP address from an incoming query and the health status of the pool or server. Here, clients have their own pool or server and are always associated with the same pool or server for the same query as long as the pool or server is green. If the health status of the pool or server turns red, NIOS switches the clients to the working pool or server and switches back when the health restores to green. For more information, see [Load Balancing Methods for DNS Traffic Control](#).

## New DNS Responses When No DNS Traffic Control Responses are Available (RFE-10212)

You now have the option to allow NIOS to either drop LBDN queries, or return DNS responses, or not return DNS responses when DNS Traffic Control responses are not available. Two new options have been introduced in the **Data Management > DNS > Grid DNS Properties/Member DNS Properties > Traffic Control** tab:

- Drop LBDN matched DNS queries during full health update: this option drops all LBDN queries when the DNS service is waiting to receive a full health status update.
- No specific behavior: this option does not return DNS responses when DNS Traffic Control responses are not available.

These options are in addition to the existing **Return DNS response if there are no DNS Traffic Control responses available** option which is selected by default. For more information, see the [Configuring DNS Traffic Control Properties](#).

## Consolidated Health Checking for DNS Traffic Control Grid Members (RFE-9427)

You can now choose the Grid members that must monitor health and share the health status. You can also select with which other members the health status is to be shared. You can do this by enabling or disabling the new **Full Health Communication** checkbox on the **Data Management > DNS > Traffic Control > Health Monitors > Advanced** tab. For more information, see the [Configuring DTC Monitors for Health Check](#).

## Notification Rule Enhancements

NIOS now includes the Delete operation type in the Outbound notification rules. The Delete operation type has been included for the DB Change DNS Record, DB Change DNS Zone, and Object Change Discovery Data event types. For more information, see the [Configuring Notification Rules](#).

## IP Address in DHCP address conflict notification (RFE-5170)

NIOS now displays the conflicting IP address along with the conflict category when an email notification is sent in case of an IPAM IP address conflict.

The content of the IB-TRAP-MIB::ibTrapDesc.0 SNMP trap is updated to STRING: DHCP address conflicts with an existing host address. [IP address].

## New Cluster Logout Event in the Syslog File (RFE-9840)

The syslog file now contains a cluster logout message to easily identify between network-related disconnects and distribution-related logouts in real time. The message is in the following format: <date:time> daemon infoblox.localdomain INFOBLOX-Grid[]: notice Cluster logout for node <node\_name>, for node clean restart.

## WAPI Performance Optimization (RFE-9986)

The performance of the WAPI GET method has been optimized for SRV, CNAME, and DNAME records.

## Grid Backup Details in the Audit Log (RFE-9614)

The audit log file now logs information about who started the database backup and where the database backup file is stored. For more information, see the [Audit Log](#).

## New CLI Command to Set DNS and Anycast Start and Restart (RFE-10176)

This release of NIOS introduces the following commands:



- `set restart_anycast_with_dns_restart` : sets DNS and anycast start and restart sequences. This command brings down the anycast service during the DNS restart or stops and redirects the traffic on the IP address of anycast to another site. You can use this command only on Grid Master.
- `show restart_anycast_with_dns_restart` : displays the status of the set `restart_anycast_with_dns_restart` command.

For more information about these commands, see the [set restart\\_anycast\\_with\\_dns\\_restart](#) and [show restart\\_anycast\\_with\\_dns\\_restart](#) topics.

## Hybrid HA Support

In NIOS 8.6, an HA setup can comprise a physical appliance and a virtual appliance. This setup is called a hybrid HA setup. For information about hybrid HA and its limitations, see the [About HA Pairs](#).

## Single Network Interface of vNIOS for GCP (RFE-9995 and RFE-9807)

This release of NIOS introduces an option to deploy vNIOS for GCP as a single network interface instance using VPC (Virtual Private Cloud) and shared VPC networks on GCP. This instance provides core network services such as DNS and IPAM services on a modular Infoblox solution. For more information, see the online Installation Guide for vNIOS for GCP at <https://infoblox-docs.atlassian.net/wiki/display/ILP/Appliances>.

## Resolving CNAME and DNAME Chains in A and AAAA Alias Records (RFE-9129)

NIOS now follows CNAME and DNAME chains if they appear as a target of an A or AAAA alias record and returns the RDATA in the final link of the CNAME and DNAME chain as the answer. The chain itself will not be present as part of the answer.

## Resetting SNMP and CLI Credentials in Network Insight (SPTYRFE-97)

If SNMP or CLI credentials become obsolete for devices polled by Network Insight, this release of NIOS introduces the following new CLI commands to reset the credentials for all affected devices at once:

- `reset snmp` : clears obsolete SNMP credentials (community strings) of devices polled by Network Insight.
- `reset cli` : clears obsolete CLI credentials (community strings) of devices polled by Network Insight.

After clearing obsolete credentials, Network Insight reguesses the credentials for each device. For information about these commands, see the [reset snmp](#) and [reset cli](#) topics.

## Credential Grouping for Discovery Devices in Network Insight

In Network Insight, you can now group credentials and assign them to devices based on their group. You can do this for devices globally, for probe members, or for individual devices.

Credentials apply to devices at the following levels:

- Grid Manager: settings apply across the Grid and all probe appliances licensed for discovery.

- Discovery probe appliances: you can use inherited Grid settings or override them.
- Individual devices: you can use inherited Grid or probe settings or override them with device-specific settings.

For more information, see the [Configuring Discovery Properties](#).

## Microsoft Server 2022 Support (RFE-12246)

NIOS 8.6 is supported on Microsoft Server 2022.

## Discovery of Cisco Viptela SDN and SD-WAN devices

You can now discover SDN and SD-WAN devices from Cisco Viptela on-premise or cloud infrastructure using Network Insight. For more information, see the [Configuring Discovery Properties](#).

## Adjustable Support Bundle Download Timeout

You can override the default timeout value for support bundle download by a custom value. For more information, see the [Downloading Support Bundles](#).

## Support for New Vendors Using Advisor

A few more new vendors can use the Advisor service to monitor their device lifecycle and vulnerabilities. For more information, see the [Monitoring Device Lifecycle and Vulnerabilities Using Advisor](#).

## Display of Source Device for Discovered Networks

You can now view the device on which a network is discovered by Network Insight. For more information, see the [Viewing Network Inventory](#).

## Unbound Upgrade

The Unbound version has been upgraded to 1.10.1.

## Enabling DDNS Updates from IPv6-Only DHCP Members (RFE-5118)

You can now enable DDNS updates from IPv6-only DHCP members.

## DHCP Fingerprint Updates

NIOS now contains new and updated DHCP fingerprints and the fingerprint configuration file has been upgraded to version 10. For details about the fingerprint format, see the [DHCP Fingerprint Detection](#).

# Support Matrix

This section lists requirements for the management system you use to access the NIOS appliance. The management system is the computer from which you configure and manage the NIOS appliance. The management system must meet the requirements listed in this topic.

## Hardware Requirements

- (Minimum) 1.4 GHz CPU with 1 GB RAM available to the product GUI, and 256 Kbps connectivity to NIOS appliance
- (Recommended) 2.0 GHz (or higher) dual core CPU with 2 GB RAM
- Monitor Resolution:
  - 1280 x 768 (minimum)
  - 1280 x 1024 or better (recommended)

## Software Requirements

For CLI access:

- Secure Socket Shell (SSH) client that supports SSHv2
- Terminal emulation program, such as minicom or Hilgraeve Hyperterminal®

For Grid Manager:

- JavaScript
- SSL version 3
- TLS version 1 connection

## Supported Browsers

For information about supported browsers, see the NIOS Release Notes on the Infoblox Support Site.

# Getting Started

This section helps you to get started using NIOS and familiarize yourself with NIOS terminology and UI conventions.






To...	See...
Understand NIOS and Infoblox terminology	<a href="#">Glossary of Terms</a>
Understand the underlying architecture of various NIOS features and components	<a href="#">SSL and TLS Protocols</a>
Know the NIOS system requirements	<a href="#">Support Matrix</a>
Get familiar with the Grid Manager interface	<a href="#">About the Grid Manager Interface</a>
Use NIOS through the CLI	<a href="#">Using the NIOS CLI</a>
Create and manage administrator, users, roles, and groups	<a href="#">Managing Administrators</a>
Learn about different types of licenses and how to add and manage licenses	<a href="#">Managing Licenses</a>
Create NIOS reports	<a href="#">Infoblox Reporting and Analytics</a>




















## Grid Manager Icons













This section contains the following information about icons used in Grid Manager, System Manager, and Orchestration Server Manager:

- **Icon:** The graphical display of an icon.
- **Icon Name:** The icon name.
- **Description:** The task that Grid Manager performs after you click the icon.
- **Tab/Table/Panel:** Lists the tab, table, or panel in which the icon appears.


The following are common icons that appear in most of the tabs, tables, panels, and in the Toolbar:

Icon	Icon Name	Description
	Active User	Indicates a user is active on the Microsoft server
	Add	Adds an object
	Add Bookmark	Adds a bookmark for an object and displays it in the Bookmarks panel
	Arrow (Down)	Moves an object down in a list
	Arrow (Up)	Moves an object up in a list














Icon	Icon Name	Description
	Clear	Clears specific data
	Clock	Displays a drop-down list for time
	Configure/Manage	Configures or manages a specific function
	Delete	Deletes an object
	Disabled	Indicates a disabled object
	Download/Import	Downloads/imports a file or data
	Edit	Displays the corresponding editor for modifying object configurations
	Execute Now	Executes a scheduled task immediately
	Export	Exports data displayed in the current panel
	Extensible Attribute	Configures extensible attributes for the selected object
	Flat View	Displays a list of objects in a flat view
	Hierarchical view	Displays objects in a hierarchical view
	Help	Displays information about an object
	Indexer	Indicates that the reporting member functions as an indexer
	Information	Displays informational data about an object
	Locked	Indicates a locked object
	Microsoft Server	Indicates a Microsoft server
	Pause	Pauses a process
	Print	Prints the information in the current panel

Icon	Icon Name	Description
	Restart	Restarts services on appliances
	Refresh	Refreshes the current page or table
	Report	Displays a report, such as the capacity report
	Search	Searches for specific objects
	Search Head	Indicates that the reporting member functions as a search head
	Start	Starts a process
	Stop	Stops a process
	Unlocked	Indicates an unlocked object
	User	Indicates a user has logged out of the Microsoft server
	User Profile	Configures a user profile
	View	Lists data in the current panel or lists detailed status about an object
	Warning	Indicates a warning message





The following icons appear in the **Data Management** tab:

Icon	Icon Name	Description	Tab/Table/Panel
	Configure	<ul style="list-style-type: none"> <li>Configures DHCP properties</li> <li>Configures File Distribution properties</li> <li>Configures Licenses</li> </ul>	<ul style="list-style-type: none"> <li><b>Data Management</b> tab -&gt; <b>DHCP</b> tab -&gt; <b>Toolbar</b></li> <li><b>Data Management</b> tab -&gt; <b>DHCP</b> tab -&gt; <b>Toolbar</b></li> <li><b>Grid</b> tab-&gt; <b>Grid Manager</b> tab -&gt; <b>Toolbar</b></li> </ul>




Icon	Icon Name	Description	Tab/Table/Panel
	Conflict	Indicates an IP address conflict	<b>Data Management</b> tab -> <b>IPAM</b> tab -> Net Map
	Convert	Converts an object	<b>Data Management</b> tab -> <b>IPAM</b> tab -> <i>network</i> -> IP Map -> Toolbar
	Discovery	Performs a network discovery	<b>Data Management</b> tab -> <b>IPAM</b> tab -> Toolbar
	Force HA Failover	Forces an HA failover	<b>Data Management</b> tab -> <b>DHCP</b> tab -> Toolbar
	Force Recovery	Forces a recovery	<b>Data Management</b> tab -> <b>DHCP</b> tab -> <b>Members</b> tab -> <b>Failover Associations</b> tab -> Toolbar
	Grid Manager	Indicates the Grid Master	<b>Data Management</b> tab -> <b>DHCP</b> tab -> <b>Members</b> tab -> <b>Data Management</b> tab -> <b>IPAM</b> tab
	Grid Manager Candidate	Indicates the Grid Master candidate	<b>Data Management</b> tab -> <b>DHCP</b> tab -> <b>Members</b> tab -> <b>Data Management</b> tab -> <b>IPAM</b> tab
	Grid Member	Indicates the Grid member	<b>Data Management</b> tab -> <b>DHCP</b> tab -> <b>Members</b> tab -> <b>Data Management</b> tab -> <b>IPAM</b> tab
	Join	Joins networks	<b>Data Management</b> tab -> <b>IPAM</b> tab -> <i>network</i> -> Toolbar
	Key-signing Key Rollover	Indicates the key-signing key that is due to rollover	<b>Data Management</b> tab -> <b>DNS</b> tab
	Microsoft Server	Indicates a Microsoft server	<b>Data Management</b> tab -> <b>DHCP</b> tab -> <b>Members</b> tab -> <b>Data Management</b> tab -> <b>IPAM</b> tab
	Multi-Ping	Pings all the addresses in a network	<b>Data Management</b> tab -> <b>IPAM</b> tab -> IP Map -> Toolbar
	Network Container	Indicates a non-cloud network container	<b>Data Management</b> tab -> <b>IPAM</b> tab or <b>DHCP</b> tab

Icon	Icon Name	Description	Tab/Table/Panel
	Network	Indicates a non-cloud network or a leaf network in a network container	<b>Data Management</b> tab -> <b>IPAM</b> tab or <b>DHCP</b> tab
	Network Container (for <b>Cloud</b> platform appliance)	Indicates a <b>cloud</b> network container	<b>Cloud</b> tab -> <b>Networks</b> tab or <b>Data Management</b> tab -> <b>IPAM</b> tab or <b>DHCP</b> tab
	Network (for <b>Cloud</b> platform appliance)	Indicates a <b>cloud</b> network	<b>Cloud</b> tab -> <b>Networks</b> tab or <b>Data Management</b> tab -> <b>IPAM</b> tab or <b>DHCP</b> tab
	Network (Disabled)	Indicates a disabled network	<b>Data Management</b> tab -> <b>IPAM</b> tab or <b>DHCP</b> tab
	Microsoft Network	Indicates a network with Microsoft servers	<b>Data Management</b> tab -> <b>IPAM</b> tab or <b>DHCP</b> tab
	Infoblox Network	Indicates a network with Infoblox appliances	<b>Data Management</b> tab -> <b>IPAM</b> tab or <b>DHCP</b> tab
	Ping	Pings an IP address	<b>Data Management</b> tab -> <b>IPAM</b> tab -> IP Map -> Toolbar
	Properties	Configures Grid DNS properties	<b>Data Management</b> tab -> <b>DNS</b> tab -> Toolbar
	Reclaim	Reclaims an IP address	<b>Data Management</b> tab -> <b>IPAM</b> tab -> IP Map -> Toolbar
	Resize	Resizes a network	<b>Data Management</b> tab -> <b>IPAM</b> tab -> <i>network</i> -> Toolbar
	Resolve Conflict	Resolves an IP address conflict	<b>Data Management</b> tab -> <b>IPAM</b> tab -> IP Map -> Toolbar
	Set Partner Down	Sets partner down	<b>Data Management</b> tab -> <b>DHCP</b> tab -> <b>Members</b> tab -> <b>Failover Associations</b> tab -> Toolbar
	Split Network	Splits a network	<b>Data Management</b> tab -> <b>IPAM</b> tab -> <i>network</i> -> Toolbar
	DNSSEC status	Displays status for DNSSEC	<b>Data Management</b> tab -> <b>DNS</b> tab -> Toolbar











Icon	Icon Name	Description	Tab/Table/Panel
	Secondary Zone Status	Displays status for the secondary zone	<b>Data Management</b> tab -> <b>DNS</b> tab
	Zoom In	Zooms in to the selected network	<b>Data Management</b> tab -> <b>IPAM</b> tab -> Net Map
	Zoom Out	Zooms out from the selected network	<b>Data Management</b> tab -> <b>IPAM</b> tab -> Net Map
	Directory	Indicates a directory	<b>Data Management</b> tab -> <b>File Distribution</b> tab

The following icons appear in the **Smart Folders** tab:




Icon	Icon Name	Description	Tab/Table/Panel
	Smart Folder	Lists a smart folder	<b>Smart Folders</b> tab
	Smart Folder (Group By)	Lists smart folders in a group-by list	<b>Smart Folders</b> tab
	Smart Folder (Link)	Indicates a link to the smart folder	<b>Smart Folders</b> tab and other selectors

The following icons appear in the **Grid** tab:





Icon	Icon Name	Description	Tab/Table/Panel
	Backup	Backs up the configuration file and database	<b>Grid</b> tab-> <b>Grid Manager</b> tab -> Toolbar
	Restore	Restores the configuration file and database	<b>Grid</b> tab-> <b>Grid Manager</b> tab -> Toolbar
	bloxTools	Performs bloxTools services	<b>Grid</b> tab-> <b>Grid Manager</b> tab -> Toolbar
	Certificate	Creates, generates, uploads, or downloads an HTTPS certificate	<b>Grid</b> tab-> <b>Grid Manager</b> tab -> <b>Members</b> tab -> <i>member</i> -> Toolbar
	Control	Restarts, reboots, or shuts down a member	<b>Grid</b> tab-> <b>Grid Manager</b> tab -> <b>Members</b> tab -> <i>member</i> -> Toolbar

Icon	Icon Name	Description	Tab/Table/Panel
	Manage Services	Manages member services	<b>Grid</b> tab-> <b>Grid Manager</b> tab -> <b>Members</b> tab -> <i>member</i>
	Syslog	Displays the syslog file	<b>Grid</b> tab-> <b>Grid Manager</b> tab -> <b>Members</b> tab -> <i>member</i> -> Toolbar
	Traffic Capture	Captures the traffic report on a member	<b>Grid</b> tab-> <b>Grid Manager</b> tab -> <b>Members</b> tab -> <i>member</i> -> Toolbar



The following icons appear in the **Administration** tab:

Icon	Icon Name	Description	Tab/Table/Panel
	Execute Now	Executes a scheduled task immediately	<b>Administration</b> tab -> Toolbar
	Overlap	Shows overlapping permissions	<b>Administration</b> tab -> <b>Permissions</b> tab
	Reschedule	Reschedules a task	Reschedules a task <b>Administration</b> tab -> Toolbar




The following icons appear in the **Finder** panel:

Icon	Icon Name	Description
	Bookmarks	Lists all bookmarked objects
	Recycle Bin	Lists all deleted objects
	Smart Folders	Lists all smart folders
	URL Links	Adds URL links

The following icons appear in the **Load Balancer** related panels:

Icon	Icon Name	Description
	Traffic Management Visualizer	Views GLB object map
	DNS View Mapping	Maps NIOS DNS view to GLB DNS view

The following icons appear in Multi-Grid Manager:

Icon	Icon Name	Description
	Apply Template	Applies templates
	Delta Viewer	Views snapshots
	External Storage	Access external storage

## Glossary of Terms

The following table provides descriptions of some key terminology used in the Infoblox products. Some terms, such as Grids and high availability, are used in different ways by other networking product vendors. The alphabetically arranged table can help you understand the terms and concepts as Infoblox uses them and as they are used in this guide.

Active Node	The NIOS appliance in an HA (high availability) pair that receives, processes, and responds to all service requests. When an HA failover occurs, the active node becomes the passive node in the HA pair.
API (Application Programming Interface)	A set of rules and specifications that software programs follow to communicate with each other. It serves as an interface between different software programs and facilitates their interaction. Infoblox provides a Perl API to help facilitate the integration of Infoblox NIOS appliances into network environments. It is an alternate method to the GUI (graphical user interface) in which you use a mouse pointer to click and select options and items to perform tasks.
Authenticated DHCP	The process of authenticating a network device before a DHCP server assigns a lease. On Infoblox appliances, you can divide a network into segments for unauthenticated, authenticated, and guest users. The Infoblox DHCP server assigns clients to the appropriate segment based on their MAC addresses and authentication credentials.
BIND (Berkeley Internet Name Domain)	The most commonly used DNS server on the Internet. It allows a standard way to name objects and resource records in distributed UNIX environments. It also provides operations to store and retrieve information of these objects and records.
bloxSYNC	An Infoblox proprietary mechanism for secure, real-time synchronization of the database that maintains the data, system configuration, and protocol service configuration between the active and passive nodes of an HA pair. With bloxSYNC, the nodes continuously synchronize changes of their configurations and states. When a failover occurs, the passive node can quickly take over services from the active node.
bloxTools	An Infoblox pre-installed environment provides tools to create custom applications that facilitate administrative tasks for an organization.
Bucket	A bucket contains indexed data.
Bulk Host	If you need to add a large number of A and PTR records, you can have the NIOS appliance add them as a group and automatically assign host names based on a range of IP addresses and the host name format you specify. Such a group of records is called a bulk host, which the appliance manages and displays as a single bulk host record.

Captive Portal	An Infoblox service that you enable on Grid members to register users, guest users, or both types of users for authentication purposes on network segments that you define using the authenticated DHCP feature.
CIDR (Classless Inter-Domain Routing) Notation	A compact specification of an IPv4 or IPv6 address and its associated routing prefix. For example, the CIDR notation of 192.168.100.1/24 represents the IPv4 address of 192.168.100.1 and its routing prefix of 192.168.100.0, or its subnet mask of 255.255.255.0. The CIDR notation of 2001:DB8::/48 represents the IPv6 addresses from 2001:DB8:0:0:0:0:0:0 to 2001:DB8:0:FFFF:FFFF:FFFF:FFFF:FFFF.
CLI (Command-line Interface)	A way to interact with Infoblox products by typing text-only commands to perform specific tasks.
Cluster Grouping	Grouping reporting appliances as a disaster recovery measure.
Dashboard	Your home page on Infoblox Multi-Grid Manager, Grid Manager, and System Manager. It provides easy access to tasks and to the status of your Grids and networks. It also provides various widgets for viewing and managing data.
DDNS (Dynamic DNS)	The automatic updating of real-time DNS configuration changes and other information on a DNS server when a network device is assigned a new IP address.
DHCP (Dynamic Host Configuration Protocol)	A configuration protocol that provides address assignments to network devices within a network. It keeps track of network configuration for each network device.
DHCP Failover Association	The pairing of two DHCP servers that establish a TCP connection for their communications. The servers form a pair of DHCP failover peers and provide DHCP protocol redundancy to minimize DHCP service outages.
DHCP Filter	A set of criteria and rules used to screen requesting hosts by matching MAC addresses, relay agent identifiers, DHCP options, or RADIUS authentication results.
DHCP Template	A set of predefined properties that you use to create IPv4 and IPv6 DHCP objects, such as networks and DHCP ranges, on the Infoblox appliance.
DIW (Data Import Wizard)	An Infoblox software tool that facilitates the import of DNS, DHCP, and TFTP data from legacy servers to Infoblox NIOS appliances. DIW supports DNS data import in the following formats: BIND 9, BIND 8, BIND 4, Microsoft DNS, Lucent VitalQIP, and Nortel NetID. It supports DHCP data import in the following formats: ISC DHCP, Microsoft DHCP, Lucent VitalQIP, and Nortel NetID.
DNS (Domain Name System)	A hierarchical naming system that translates domain names of any network devices into IP addresses for the purpose of locating and addressing these devices worldwide.
DNS View	On Infoblox appliances, a DNS view provides the ability to serve one version of DNS data to one set of clients and another version to another set of clients. With DNS views, the Infoblox appliance can provide a different answer to the same DNS query, depending on the source and match destinations of the query.
DNSSEC (Domain Name System Security Extensions)	A suite of IETF (Internet Engineering Task Force) specifications to secure certain kinds of information provided by DNS for use on IP networks. It is a set of extensions to DNS, which provide DNS resolvers with the original authentication of DNS data, authenticated denial of existence, and data integrity.
DNSone™	The software package that enables Infoblox appliances to provide DNS, DHCP and TFTP services. You can add the Grid upgrade to Infoblox appliances running DNSone.

Endpoint	An IP device such as a personal computer, laptop, or mobile handheld device. This term is often used in a security context.
Extensible Attribute	Metadata you define to capture additional information about an object managed by the Infoblox NIOS appliance. You can use predefined attributes or create your own. You can also specify required attributes and restrict the values that users can enter for each attribute.
Filters	Criteria the Infoblox NIOS appliance uses to request specific information in the database. You can use filters to control the amount and the kind of data displayed in a panel or table in Infoblox Multi-Grid Manager, Grid Manager, and System Manager.
FQDN (fully qualified domain name)	A complete domain name that specifies its exact location in the hierarchy of the DNS. It specifies all the domain levels, including the top-level domain and the root domain.
FTP (File Transfer Protocol)	A standard network protocol used to transfer files from one network device to another over a TCP-based network, such as the Internet. FTP is built on a client-server architecture and utilizes separate control and data connections between the client and server.
Gateway	The default router for the immediate network segment of an interface.
Grid™ Technology	Infoblox's unique and patented high availability Grid technology ensures network reliability. The Infoblox Grid provides resilient network services, failover, recovery, and seamless maintenance for an Infoblox deployment inside a single building, across a networked campus, or between remote locations. The Infoblox Grid establishes a distributed relationship between individual or paired appliances to remove single points of failure and other operational risks inherent in legacy DNS, DHCP, and IP address management infrastructure.
Grid Manager	The NIOS web interface that provides access to your Grid for performing IPAM, DNS, and DHCP management and other administration tasks.
Grid Master	The Grid member in an Infoblox Grid that maintains the NIOS database that is distributed among all members of the Grid. You connect to the Grid Master to configure and monitor the entire Grid.
Grid Member	Any single Infoblox NIOS appliance or HA pair that belongs to a Grid. Each member can use the data and services of the Grid. You can also modify settings so that a Grid member can use unique data and member-specific services.
HA Pair	Two physical Infoblox NIOS appliances that are linked to perform as a single virtual appliance in an HA (high availability) configuration. The HA configuration provides hardware redundancy to minimize service outages. In this configuration, one appliance is the active node and the other is the passive node.
Host Record	On Infoblox appliances, host records provide a unique approach that enables you to manage multiple DNS records and DHCP and IPAM data collectively, as one object on the appliance.
IBOS (Infoblox Orchestration Server)	IBOS is the Infoblox IF-MAP (Interface to Metadata Access Points) server that contains a searchable database for storing state information about network resources. It is the central point with which IF-MAP clients communicate to send and retrieve real-time information defined in the IF-MAP data format.
IF-MAP (Interface for Metadata Access Points)	An open standard client-server protocol developed by the Trusted Computing Group as one of the core protocols of the TNC (Trusted Network Connect) open architecture. IF-MAP allows network resources to share real-time information.

IP Map	In Infoblox Grid Manager or System Manager, this is a graphical representation of all IPv4 addresses in a given subnet.
IPAM (IP Address Management)	Infoblox IPAM provides a means of planning, tracking, and managing IP address space in a network. It glues DNS and DHCP services together so that each service is aware of changes in the other. The Infoblox IPAM implementation offers an IP address-centric approach so you can manage your networks and IP addresses through a centralized GUI.
Leaf Network	On Infoblox appliances, a network that does not contain any subnets. Lease Logging Member An Infoblox Grid member that is designated to collect DHCP lease events.
License Pool	A license pool is a container associated with the Grid and holds dynamic licenses for a specific feature. Licenses in the pool can be dynamically allocated to and deallocated from Grid members. When not in use, dynamic licenses are released back to the pool for future allocation. There is no expiration for dynamic licenses.
Limited-Access User	An admin user account that has specific roles and permissions assigned. Limited-access users have restricted access to Infoblox Multi-Grid Manager, Grid Manager, and System Manager, and can only perform certain tasks based on their assigned roles and permissions.
Lite Upgrade	On Infoblox appliances, a lite upgrade occurs when there are incremental changes to the NIOS software that do not require any change to the database. The appliance can perform a lite upgrade only if the format of the database between the existing NIOS version and the upgrade version is the same. In general, when you upgrade from a major release to a patch release or a patch release to another patch release, you are performing a lite upgrade.
Loopback Interface	On Infoblox appliances, the virtual network interface on which you can consolidate DNS servers for migration purposes, add anycast addresses to improve the performance of the DNS service, and separate DNS traffic.
Managing Member	An Infoblox Grid member that is configured to manage Microsoft DNS and DHCP servers.
Master Candidate	An Infoblox Grid member that is designated to assume the role of the Grid Master as a disaster recovery measure.
Master Grid	A group of Infoblox appliances that are connected to provide a single point of administration for multiple Grids and network management of these Grids.
Master Grid Member	Any single Infoblox appliance or HA pair that belongs to the Master Grid. All Master Grid members serve as Master Candidates.
Multi-Grid Manager	The NIOS web interface that provides access to the Master Grid, from which you can manage multiple Grids and their networks.
Multi-Grid Master	The Infoblox Master Grid member that maintains the NIOS database that is distributed among all Master Grid members. You connect to Multi-Grid Manager to configure and monitor the Master Grid.
Multi-Grid Master Candidate	An Infoblox Master Grid member that is designated to assume the role of the Multi-Grid Master as a disaster recovery measure.
Name Server Group	On Infoblox appliances, a server group that contains one primary DNS server and/or one or more secondary DNS servers. Specifying a single name server group can simplify DNS zone creation.

NAT (Network Address Translation) Group	A group of Infoblox Grid members that are configured on the same side of a NAT appliance. In a Grid configuration where the Grid Master is configured behind a NAT appliance and there are Grid members on both sides of the NAT appliance, it is necessary to create a NAT group to ensure that the Grid Master and Grid members use the correct NAT and interface addresses for Grid communications.
Network Block	On Infoblox appliances, an IP address space that is defined in the Master Grid. A network block can consist of other network blocks, network containers, and leaf networks.
Network Container	On Infoblox appliances, an automatically created container of multiple networks that are subnets of the IP address space configured for the network container. A network container cannot be assigned to a Grid member or be directly created.
Network Discovery	A set of tools provided by the Infoblox NIOS appliance for detecting active hosts on specified networks and specified VMware vSphere servers.
Network Map	In Infoblox Grid Manager and System Manager, Network Map presents a complete view of your network space, including the different types of networks that are in it and its unused address space. You can use Network Map to design and plan your network infrastructure, configure and manage individual networks, and evaluate their utilization.
Network Mask or Netmask	A numeric representation of the bits that are used to split an IP address into the network portion and the host portion. In Infoblox products, this is represented by either quad-dotted decimal representation or CIDR notation for IPv4 network masks, or by CIDR notation for IPv6 network masks.
Network View	On Infoblox appliances, a single routing domain with its own networks and shared networks. A network view can contain both IPv4 and IPv6 networks. All networks must belong to a network view on the Infoblox appliance.
NIOS	An Infoblox proprietary system that powers Infoblox solutions with an embedded processor that delivers core network services. It is the operating system that runs on the NIOS appliances—a security-hardened, real-time set of appliances built to ensure the non-stop operation of network infrastructure. NIOS automates the error-prone and time-consuming manual tasks associated with deploying and managing IPAM, DNS, and DHCP required for continuous IP network availability and business uptime.
NIOS Virtual Appliance	Any Infoblox supported platform, such as AWS, VMWare, Azure appliances that runs the vNIOS software. These appliances are also known as the vNIOS appliances.
Node	A single Infoblox appliance of an HA (high availability) pair. An HA pair consists of an active node and a passive node.
NTP (Network Time Protocol)	A protocol for synchronizing the clocks of computer systems over packet-switched, variable latency data networks; it essentially keeps network devices on a common clock by resisting the effects of variable latency by means of a jitter buffer.
Passive Node	The Infoblox NIOS appliance in an HA pair that constantly keeps its database synchronized with that of the active node, so it can take over core network services when an HA failover occurs. When an HA failover occurs, the passive node becomes the active node in the HA pair.
PortIQ	An Infoblox switch port appliance that enables quick discovery of the Ethernet switch ports. PortIQ identifies ports that are not fully utilized and those that exceed their capacity. You can use PortIQ to troubleshoot LAN environments.

Quick Filter	A filter that stores specific filter criteria for requesting information displayed in a specific panel in Infoblox Multi-Grid Manager, Grid Manager, and System Manager. For more information, see "Filter."
Overlapping Network	On Infoblox appliances, a network that exists in multiple locations, which can be multiple Grids in the Master Grid or within various network views in a Grid.
Replication	Database distribution among the Infoblox Grid Master and Grid members as well as among the Multi-Grid Master and Master Grid members.
Replication Factor (Reporting - Multi-Site Cluster)	The number of copies of reporting data in each bucket that the cluster maintains.
Reservation	On Infoblox appliances, a static IP address that you create for future use. A reservation is a pre-provisioned fixed address. You can reserve this static IP address on the NIOS appliance and assign it to a client in the future.
Resource Records	A collection of data in the DNS server database. Each resource record specifies information about a DNS object. For example, an A (address mapping) record maps a host name to an IP address, and a PTR (reverse-lookup pointer) record maps an IP address to a host name. The DNS server uses these records to answer queries.
Roaming Host	On Infoblox appliances, a host with a dynamically assigned IP address and a specific set of properties and DHCP options. When you create a roaming host for a network device, the device can receive any dynamically assigned address from the network to which it belongs.
Scope	A DHCP address range on a Microsoft server. Microsoft scope information is converted to equivalent DHCP range information after Microsoft data is synchronized with the NIOS appliance.
Search Factor	The number of searchable copies of reporting data in each bucket that the cluster maintains.
Shared Network	On Infoblox appliances, a network segment to which you assign two or more subnets. When subnets in a shared network contain IP addresses that are available for dynamic allocation, the addresses are put into a common pool for allocation when client requests arise.
Shared Record Group	On Infoblox appliances, a set of resource records that you add to multiple DNS zones. You can create resource records in a group and share the group among multiple zones. The zones handle the shared resource records as any other resource record.
SSO (Single Sign On)	An Infoblox feature that allows you to automatically sign in to selected Grids from the Master Grid, without having to log in to each individual Grid each time you sign on.
Smart Folder	On Infoblox appliances, a virtual folder in which you place the results of filter criteria that you select to request specific data in the NIOS database. Once you set up a smart folder, the appliance displays up-to-date information based on your filter and grouping criteria each time you access the folder.
Subnet (or network)	A logical division of an IP network. A subnet of network may also be called a network. For example, 10.1.0.0/16 is a subnet of 10.0.0.0/8, and fc80:8:8:16::/64 is a subnet of fc80:8:8::/48.
Superscope	On a Microsoft server, superscope comprises multiple scopes or DHCP address ranges created on a single physical network segment. Microsoft superscope information is converted to equivalent network information after Microsoft data is synchronized with the NIOS appliance.



Superuser	An admin user account that has unrestricted access to Infoblox Multi-Grid Manager, Grid Manager, or System Manager.
Support Bundle	A tar.gz file that contains configuration files and system files of the Infoblox NIOS appliance. You can download a support bundle for an independent appliance and for each member in a Grid.
System Manager	The NIOS web interface that provides access to an independent appliance (single or HA) for performing IPAM, DNS, and DHCP management and other administration tasks.
TFTP (Trivial File Transfer Protocol)	A data transfer service that provides devices—such as phones, RFID readers, IP cameras, and other devices—with up-to-date software and configuration data.
Traffic Capture	An Infoblox tool that captures the traffic on one or all of the ports on a NIOS appliance. The NIOS appliance saves all captured traffic in a .cap file and compresses it into a .tar.gz file.
Upgrade Group	On Infoblox appliances, a group of Grid members that you put together so you can perform software distribution and upgrade at the same time.
VIP (Virtual IP)	On Infoblox appliances, the shared IP address of an HA pair. A VIP address links to the HA port on the active node of an HA pair.
VRID (Virtual Router ID)	VRID identifies the VRRP (Virtual Router Redundancy Protocol) HA pair to which the Infoblox appliance belongs. Through VRID, two HA nodes identify each other as belonging to the same HA pair, and they obtain a virtual MAC address to share with a VIP. A VRID can be any number between 1 and 255, and it must be unique on the local LAN so that it does not conflict with any other Infoblox appliances using VRRP on the same subnet.
vNIOS	The virtual version of NIOS. You can install Infoblox vNIOS software on any supported virtual platform and configure the system as a vNIOS virtual appliance.
VRRP (Virtual Router Redundancy Protocol)	An industry standard MAC address level HA failover mechanism.

# Installing NIOS

NIOS comes with its own hardware box (appliance) on which NIOS will be installed. Ensure that you verify the hardware that you receive.

## Verifying the Hardware

To verify the secure delivery of the hardware:

- Use the tracking number of the order to review the status of the shipment.
- Inspect the tamper-evident seals for any signs of tampering.
- Verify the product by comparing the shipping slip with the invoice.

## Prerequisites

Before you begin the configuration, ensure that you have all the necessary components. The following are needed and must be acquired before continuing with this guidance:

- Supported NIOS appliances
- A management station or computer from which you configure and manage the NIOS appliance. See [Support Matrix](#) for the system and browser requirements.
- The IP address of the appliance on your network.

## Supported Appliances

For the list of supported physical appliances and virtual appliances that NIOS is supported on, see the Release Notes.

## Downloading the NIOS Software

You can download NIOS from the Infoblox Support web site at <https://support.infoblox.com/>. You must be a registered user to access this web site.

To download the software:

1. Log on to the Infoblox Support web site at <https://support.infoblox.com/>
2. Click the **Downloads** tab.
3. Select **NIOS/vNIOS** from the **Infoblox Software** drop-down list.
4. Select the **General maintenance products with full engineering support for routine patches and bug fixes on all significant issues** option.
5. Select **NIOS 8.6.1** from the **Select version** drop-down list.

The NIOS installer files are available for download. The installer files are of two sizes:

- Files of size 250 GB
- Files of size 70 GB

Infoblox recommends using a minimum size of 70 GB for any of these files that has resizable as part of the file name and you can resize them depending on your requirement and deployment.

If you download the resizable files, a default disk size is defined on the VM. You then change the default disk size using the VMware UI. When you start the VM for the first time, NIOS detects how much disk space has been reserved for the virtual disk and allocates partitions within that space. You can increase the disk size upto a maximum of 2.5 terabytes.

If you want to resize the disk space from the current 250 GB to a lower size, remove the member from the Grid and add a replacement member with the required disk size. If you want to resize a smaller disk to a higher size, back up the Grid database, and deploy the new instance with the required disk size. After the deployment, restore the backed up

database. Infoblox recommends configuring a scheduled backup so that if the disk limit is reached and the system is unresponsive, the earlier backup can be restored.

Infoblox recommends configuring instances of only 250 GB and higher as Grid Master and Grid Master Candidate.

 **Note**

- Images of types .vhd and .raw comprise only one file each and can be resized. The minimum size of the .vhd image is 250 GB and that of the .raw image is 43 GB.
- Infoblox recommends that you provision 70 GB or more disk space for the NIOS instance while deploying resizable images.
- If you add members to a FIPS mode -enabled Grid Master, ensure that the members have the same NIOS build version. That is, the same NIOS build must be installed on Grid Master and members before joining. You can perform upgrades and install hotfixes in FIPS mode.

## Limitations of Using the Scalable Image File

The following are the limitations of using the scalable NIOS image:

- You can use the scalable image only on the following appliances: IB-V805, IB-V1405, IB-V2205, IB-V4005, IB-V815, IB-V825, IB-V1415, IB-V1425, IB-V2215, IB-V2225, IB-V4015, IB-V4025, IB-FLEX, CP-V805, CP-V1405, CP-V2205.
- If you use ovftool to deploy the OVA image, there is no command line option to increase the disk-size. As a workaround, you can use ovftool to upload the image but before powering on the VM, you will need to use the UI to change the disk size based on your requirement.
- The Parental Control feature fails to start because of no space left on the device when downloading Zvelo data.
- The primary disk size is full if you enable the bloxTools service on the member provisioned with the lower disk size.
- Infoblox recommends to not run DNS/DHCP/updates on the member provisioned with the lower disk size.
- Do not try to join a resizable disk image to a Grid running an NIOS version earlier than 8.5.
- Infoblox does not recommend auto-synchronizing or downgrading resizable disk images to NIOS versions earlier than 8.5.
- If you have resizable members whose provisioned disk space is lower than 85 GB, disable the rsync-batch mode by running the `set upgrade_dist rsync_batch disable` command on Grid Master.
- If you have a reporting member as part of your Grid, then you must have a minimum disk space of 250 GB by default.

# Upgrading NIOS

This topic explains how to manage upgrade groups and perform software upgrades and downgrades for NIOS appliances. It also describes how to back up and restore configuration files.

It includes the following sections:

- [Upgrading NIOS Software](#)
- [Lite Upgrades](#)
- [Viewing Software Versions](#)
- [Downgrading Software](#)
- [Reverting the Grid to the Previously Running Software](#)
- [Applying Hotfixes](#)
- [Using Database Snapshots](#)
- [Downloading Support Bundles](#)
- [Guidelines for Upgrading, Backing Up, and Restoring Data](#)
- [Guidelines for Upgrading the Reporting and Analytics Solution](#)
- [Guidelines for Scheduling Full Upgrades](#)
- [Backing Up and Restoring Configuration Files](#)
- [Managing Upgrade Groups](#)
- [Full Upgrades](#)

## Upgrading NIOS Software

Infoblox frequently releases updated NIOS software. Contact Infoblox Technical Support to learn which file name to use when downloading a new upgrade file, or watch your email for periodic notifications that a new software upgrade is available. To get the latest upgrade, your local network must be capable of downloading a file from the Internet. After you download and store the new upgrade file on your local network, complete the following tasks to upgrade an Infoblox independent appliance or a Grid.

- Upload the new software to the Grid Master, as described in [Uploading NIOS Software](#) below.
- Distribute the software upgrade files, as described in [Distributing Software Upgrade Files](#) below.
- Optionally, test the upgrade, as described in [Testing Software Upgrades](#) below.
- Perform the software upgrade, as described in [Performing Software Upgrades](#) below.

Before upgrading, Infoblox recommends that all members in the Grid be connected to the network and operating normally. If one or more members are offline when you upgrade the Grid, they automatically receive the distributed software and upgrade when they join the Grid or come back online.



### Note

- You cannot upgrade directly to NIOS 5.x from NIOS releases earlier than 4.2r4. Refer to the release notes for the appropriate upgrade and revert paths.
- In a Grid that is configured with an HA pair, Infoblox does not recommend that you disconnect any node of the HA Grid Master and then join it back after installing a NIOS version that does not match with the current running version of NIOS on the other node. If an upgrade operation is performed this way, the Grid displays an unexpected behavior and will need a manual intervention of the Support team to recover the Grid.
- The shared secret that you enter when adding a RADIUS authentication server in the *Add RADIUS Authentication Service* wizard > **RADIUS Servers** > **Shared Secret** field must be between 4 and 64 characters (inclusive) in length. Otherwise, the upgrade will fail.



### Caution

Do not attempt to add or remove a member, or convert an HA pair to single members or vice versa during a distribution or upgrade.

When you upgrade from NIOS 6.4.0 to a later release, you can start, stop, or restart DNS and DHCP services, or only the DHCP service on a member that has not been upgraded. When you start, stop, or restart other services, such as reporting or file distribution, the operation is put in queue for execution until after the targeted member has been upgraded.

## Uploading NIOS Software

After you download the NIOS software upgrade to your management station, upload it to the Grid Master, as follows:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Upload** in the panel or from the Toolbar.
2. Navigate to the directory where you have stored the NIOS software upgrade, and then click **Open** or **Upload**.

The appliance uploads the file and displays the status of the upload in the status bar. You can click the Stop icon in the status bar to stop the upload. Ensure that you do not navigate away from the **Upgrade** tab until after the upload is complete. Otherwise, the upload process stops.



### Note

When you upload the NIOS software upgrade to an HA Grid Master, only the active node receives the software. The passive node does not. Therefore, if the Grid Master fails over before a distribution starts, you must upload the software again. If you do not, the distribution fails because the new active node does not have the uploaded software.

## Distributing Software Upgrade Files

Distributing the software upgrade files involves unpacking the software files and loading the new software. When you perform a distribution, the NIOS appliance loads the new software code into an alternate disk partition, which overwrites any previously saved version of code that is already there. Therefore, starting the distribution disables the appliance from reverting to a release prior to the current version.

The time this process takes depends on the number of appliances to which the software is distributed; the more appliances, the longer it takes. Therefore, you might want to schedule the Grid distribution during times when your network is less busy. You can distribute the software immediately or schedule the distribution of any software upgrade file, even if it is not Upgrade Lite compatible.

### Distributing Software Immediately

The Grid Master distributes the software upgrade to each member in the Grid, including itself. As an alternative to scheduling the Grid distribution (see Scheduling Distributions below), you can distribute the software upgrade throughout the Grid immediately, as follows:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Distribute** -> **Distribute Now** from the Toolbar.
2. In the confirmation dialog box, click **Yes** to start the distribution.

The distribution starts and if there is an active distribution scheduled, the appliance changes its status to inactive. The appliance distributes the upgrade files and displays the status of the distribution in the status bar. You can pause, resume, or stop the distribution by clicking the corresponding icon in the status bar.

Note that starting a manual distribution cancels a scheduled distribution.

### Scheduling Distributions

When you schedule a distribution, you schedule the distribution of the Grid Master as well as the upgrade groups, including the Default group. The Grid Master distribution must always occur before the distribution of the upgrade groups. To schedule a software distribution:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Distribute** -> **Schedule Distribution** from the Toolbar.
2. In the *Schedule Distribution* editor, complete the following:

- **Activate Distribution Schedule:** Select this to enable the distribution schedule. Clear this if you are creating a distribution schedule you plan to activate at a later date. You can configure and save information in this editor even when you deactivate a scheduled distribution.
  - **Grid Master Distribution Start Information:** Enter a Grid Master distribution date, time, and time zone. The distribution date and time must be before those of the upgrade groups.
    - **Date:** Enter a start date of the Grid Master distribution in YYYY-MM-DD (year-month-day) format. You can click the calendar icon to select a date from the calendar widget.
    - **Time:** Enter a start time of the Grid Master distribution in hh:mm:ss AM/PM (hour:minute:second in AM or PM) format. You can also select a time from the drop-down list.
    - **Time Zone:** Select a time zone that applies to the start time you enter. If this time zone is different from the Grid time zone, the appliance converts the time you enter here based on the Grid time zone, after you save this schedule. When you display this schedule again, it displays the converted time. Selecting the time zone here does not affect any time zone settings in the Grid. (For information about selecting the Grid and member time zones, see [Managing Time settings.](#))
    - **Admin Local Time:** Displays the Grid Master distribution start date and time in the time zone of the administrator, as explained in [Creating Local Admins.](#)
  - In the upgrade group table, specify the following for each upgrade group by clicking the corresponding field in each row:
    - **Start Distribution:** Specify when the distribution occurs. Select one of the following from the drop-down list:
      - **Date/Time:** Select this to configure the distribution start date, time, and time zone.
      - **After <group> :** Select **After Grid Master** to start the distribution immediately after the completion of the Grid Master distribution. Select an upgrade group that must complete its distribution before the group you are configuring. When you select this option, you cannot enter a date, time, and time zone.  
**Date, Time, and Time Zone** are enabled only when you select **Date/Time** for **Start Distribution.**
      - **Date:** Enter a distribution start date in YYYY-MM-DD (year-month-day) format. You can click the calendar icon to select a date from the calendar widget.
      - **Time:** Enter a distribution start time in hh:mm:ss AM/PM (hour:minute:second in AM or PM) format. You can select a time from the drop-down list.
      - **Time Zone:** By default, the appliance displays the time zone of the first Grid member in the Upgrade Group. You can change this time zone if you want to enter the time using a different time zone. After you save the schedule though, the appliance converts the time you entered to the time zone of the upgrade group, if it is different. (For information about setting the Grid and member time zones, see [Managing Time Settings.](#)) To change the default time zone of the upgrade group, change the time zone of the first group member, as explained in Adding Upgrade Groups, see [Managing Upgrade Groups.](#)
      - **Admin Local Time:** Displays the start date and time in the time zone of the administrator, as explained in [Creating Local Admins.](#)
      - **Distribute to Members:** Indicates whether the distribution within the group occurs simultaneously or sequentially. You cannot edit this field here. You define this when you create the upgrade group. To change this setting, see Modifying Upgrade Groups in [Managing Upgrade Groups.](#)
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Grid Manager confirms that the schedule is saved and indicates whether the distribution schedule is active. You can click the Refresh icon to refresh the information in this panel.

Note that the appliance does not save the schedule and displays an error message if the schedule contains the following:

- Circular dependencies between upgrade groups. For example, the distribution of Group A is scheduled after Group B, and the distribution of Group B is scheduled after Group A.
- The distribution time is in the past.

## Software Distribution Process

The following series of events occur after a Grid distribution starts:

- The appliance checks if a NIOS software upgrade was uploaded.
  - If the upgrade files are not uploaded, the distribution stops. The appliance displays an error message and if the distribution is scheduled, the appliance deactivates the distribution schedule.
  - If the upgrade files are uploaded, the distribution proceeds.
- A single Grid Master uploads the file to a backup partition and unpacks the contents, which overwrites any existing backup software that might have been there. For an HA Grid Master, it is the active node that uploads the file to a backup partition and unpacks the contents.
  - The Grid Master (or active node of the HA Grid Master) sends a command to all nodes that are online to copy their database and software to a backup software partition.
  - For an HA Grid Master, the active node sends the command to the passive node as well.
  - The nodes perform resynchronization on their backup partition, retrieving only the changed files from the Grid Master.
  - After the active node of an HA member receives the software, it then distributes it to the passive node.

When the distribution successfully completes, the appliance updates the distribution status and sets the schedule, if configured, to inactive. The new software is now staged on all member appliances and is ready for use. Grid Manager displays the software version in the **Distribution** field in the Grid Version Information section.

## Managing Distributions

After you start a distribution, you can pause, resume, or stop it. For information, see [Pausing and Resuming Distributions](#) and [Stopping Distributions](#) below. Grid Manager displays the status of the overall distribution as well as the status of individual members. You can view this information in the **Upgrade** tab.

### Pausing and Resuming Distributions

The following are some operational guidelines for performing a distribution:

- You cannot create new upgrade groups, add members to a group, or remove members from a group after a distribution starts.
- You can skip a member that is currently offline from a distribution. When both nodes of an HA pair are online, the skip member function is not available.

To pause a distribution:

1. From the Grid Distribution Status bar, click the Pause icon.
2. When the appliance displays a confirmation dialog box, click **Yes** to pause the distribution.

The Grid Distribution Status bar indicates the distribution is paused. For information about the distribution status of each member, see [Monitoring Distribution and Upgrade Status](#) below.

To skip a member from a distribution:

1. From the **Grid** tab, click the **Upgrade** tab, and then click **Toggle Member List View**.
2. Select a member checkbox, and then click **Skip Member** from the Toolbar. Grid Manager automatically skips the distribution of software to the members that are offline.

To resume a distribution:

1. From the Grid Distribution Status bar, click the Resume icon.
2. When the appliance displays a dialog box confirming that you want to resume the distribution, click **Yes** to continue.

Members that have not completed or started distributions that were scheduled at an earlier time resume the distribution.

### Stopping Distributions

You can stop a distribution immediately, for example, if there are offline members and you do not want to wait for them to come back online, or if you realize that you have uploaded the wrong software version. When you stop a distribution, you can do the following:

- If the Grid Master has completed its distribution, you can upgrade the Grid immediately. This forces members that do not have a complete distribution to synchronize their releases with the Grid Master.
- If the Grid Master does not have a valid distribution, you can restart the distribution.
- Upload another software upgrade.

Ending a distribution does not affect the upgrade schedule, if configured. The Grid upgrade starts as scheduled, as long as the Grid Master completes its distribution.

To stop a distribution:

1. From the Grid Distribution Status bar, click the Stop icon.
2. When the appliance displays a dialog box confirming that you want to stop the distribution, click **Yes** to continue.

## Testing Software Upgrades

After you successfully distribute a software upgrade to the Grid Master, you can test an upgrade on the Grid Master before actually implementing it. This allows you to resolve potential data migration issues before the actual upgrade. The length of time the upgrade test takes depends on the amount of data and the difference between the current NIOS version and the software upgrade. The test does not affect NIOS services and you can perform other administrative tasks during the upgrade test.

To start an upgrade test:

- From the **Grid** tab, select the **Upgrade** tab, and then click **Test Upgrade** from the Toolbar. Test upgrade is enabled only for a major upgrade (not an Upgrade Lite compatible upgrade).

After you start an upgrade test, you can view its status in the status bar. You can also stop it at any time. To stop an upgrade test:

- From the *Grid Upgrade Test Status* bar, click the Stop icon.

Note that if an admin restarts the Grid services or reboots the Grid Master, or if an HA failover occurs on the Grid Master during the upgrade test, the appliance automatically stops the test. The appliance always resets the status of the Grid to "Distributed" when it stops the upgrade test.

If the appliance encounters an error during the test, it stops the test and displays a message in the *Upgrade Status* panel indicating that the upgrade test failed and the reason for the failure, such as a data translation error or data import error. You can review the syslog for specific error messages before downloading the Support Bundle and contacting Infoblox Technical Support.

After the test successfully finishes, the appliance displays a message confirming that the test upgrade is complete.

## Performing Software Upgrades

Performing a software upgrade involves rebooting the appliances and then running the new software. Essentially, each appliance switches between the two software partitions on its system, activating the staged software and saving the previously active software and database as backup.



### Note

Before you upgrade the software, Infoblox recommends that you back up the current configuration and database. For information, see [Backing Up and Restoring Configuration Files](#).

Depending on your upgrade paths, you can upgrade to a new release immediately or you can schedule the upgrade. For information about how to upgrade immediately, see [Upgrading the Grid Immediately](#) below. Before you schedule an upgrade, ensure that you understand the limitations, as described in [Managing Upgrade Groups](#). For information about how to schedule an upgrade, see [Scheduling Upgrades](#) below.



## Upgrading the Grid Immediately

For unschedulable full upgrades, all the Grid members in the Grid must upgrade at the same time. For lite upgrades and schedulable full upgrades, you can schedule the upgrades as described in [Scheduling Upgrades](#), or you can upgrade all the Grid members at the same time.

To upgrade a Grid immediately:

- From the **Grid** tab, select the **Upgrade** tab, and then click **Upgrade** -> **Upgrade Now** from the Toolbar.



### Note

The Grid upgrades immediately and if there is an active upgrade schedule, it becomes inactive.

## Scheduling upgrades

You can schedule lite upgrades and full upgrades for certain NIOS versions. For limitations about scheduling a full upgrade, see [Managing Upgrade Groups](#). When you schedule an upgrade, you schedule the upgrade for the Grid Master and the upgrade groups, including the Default group. The Grid Master must always upgrade before the upgrade groups. Depending on your upgrade paths, you can schedule the upgrade for the Grid Master and upgrade groups at different times over a period of nine days. If you schedule an upgrade that takes more than nine days, the appliance displays a warning.

To schedule an upgrade:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Upgrade** -> **Schedule Upgrade** from the Toolbar.
2. In the *Upgrade Schedule* editor, complete the following:
  - **Activate Upgrade Schedule:** Select this to enable the upgrade schedule. Clear it if you are creating an upgrade schedule that you plan to activate at a later date. You can configure and save information in this editor even when you deactivate a distribution.
  - **Grid Master Upgrade Start Information:** Enter a Grid Master upgrade date, time, and time zone. The date and time must be before those of the upgrade groups.
    - **Date:** Enter a start date of the Grid Master upgrade in YYYY-MM-DD (year-month-day) format. You can click the calendar icon to select a date from the calendar widget.
    - **Time:** Enter a start time of the Grid Master upgrade in hh:mm:ss AM/PM (hour:minute:second in AM or PM) format. You can select a time from the drop-down list.
    - **Time Zone:** Select a time zone that applies to the start time you enter. If this time zone is different from the Grid time zone, the appliance converts the time you enter here based on the Grid time zone, after you save this schedule. When you display this schedule again, it displays the converted time. Selecting the time zone here does not affect any time zone settings in the Grid. (For information about setting the Grid and member time zones, see [Managing Time Settings](#).)
    - **Admin Local Time:** Displays the Grid Master upgrade date and start time in the time zone of the administrator, as explained in [Creating Local Admins](#).
  - In the upgrade member table, specify the following by clicking the corresponding field in each row:
    - **Group:** The name of the upgrade group. You can assign a different upgrade group by selecting the group from the drop-down list.
    - **Group Members:** When you expand an upgrade group, this field displays the group members.
    - **Warning:** This field turns yellow when there is a conflict among the upgrade groups. Hover your mouse over the field and the tooltip displays the member that contains the conflict. It also displays recommended upgrade groups in the **Group** column so you can change the group assignment to resolve the conflict. The tooltip can display one of the following: **GMC**, **DNS Primary**, **DHCP Logging Member**, or **DHCP Failover**. For information about how to resolve a conflict, see [Resolving Upgrade Warnings](#) below. Select an upgrade group from the drop-down list in the Group column to assign a different upgrade group. Click **Validate and Refresh** to validate the new group assignment.
    - **Start Upgrade:** Specify when the upgrade occurs. Select one of the following from the drop-down list:
      - **Date/Time:** Select this to configure the upgrade start date, time, and time zone.

- **After <group>** : Select **After Grid Master** to start the distribution immediately after the completion of the Grid Master distribution. Select an upgrade group that must complete its distribution before the group you are configuring. If you select this option, you cannot enter a date, time, and time zone.

**Date**, **Time**, and **Time Zone** are enabled only when you select **Date/Time** for **Start Upgrade**.

- **Date**: Enter an upgrade start date in YYYY-MM-DD (year-month-day) format. You can click the calendar icon to select a date from the calendar widget.
- **Time**: Enter an upgrade start time in hh:mm:ss AM/PM (hour:minute:second in AM or PM) format. You can select a time from the drop-down list.
- **Time Zone**: By default, the appliance displays the time zone of the first Grid member in the Upgrade Group. You can change this time zone, if you want to enter the time using a different time zone. After you save the schedule though, the appliance converts the time you entered to the time zone of the upgrade group, if it is different. (For information about setting the Grid and member time zones, see [Managing Time Settings](#).) To change the default time zone of an upgrade group, change the first group member in the Upgrade Group list, as explained in Adding Upgrade Groups, see [Modifying Upgrade Groups](#).
- **Admin Local Time**: Displays the data and time in the time zone of the administrator, as explained in [Modifying Upgrade Groups](#).
- **Upgrade Members**: Indicates whether the upgrade within the group occurs simultaneously or sequentially. You cannot edit this field here. You define this when you create the upgrade group. To change this setting see as described in Modifying Upgrade Groups, see [Modifying Upgrade Groups](#).

3. Save the configuration.

The appliance does not save the schedule and displays an error message if the schedule contains the following:

- Circular dependencies between upgrade groups; for example, the upgrade of Group A is scheduled after Group B, and the upgrade of Group B is scheduled after Group A.
- The upgrade time is in the past.

The appliance also does not save the schedule and displays a warning when there is a group assignment conflict. Otherwise, the appliance confirms that the schedule is saved and indicates whether the upgrade schedule is active.

## Resolving Upgrade Warnings

The appliance can generate the following warnings when you schedule an upgrade:

- **GMC**: To resolve this warning, put all Grid Master candidates in the first upgrade group.
- **DNS Primary**: To resolve this warning, put all the members that are serving as DNS primaries in the first upgrade group.
- **DHCP Logging Member**: To resolve this warning, put the DHCP logging member in the first upgrade group.
- **DHCP Failover**: To resolve this warning, place the peers of a DHCP failover association in separate upgrade groups. Ensure that you schedule upgrades of the failover peers close to each other to minimize configuration rules. NIOS does not allow DHCP configuration changes that affect the communication between the peers until both peers are upgraded.

## Upgrading Groups Immediately

After you schedule an upgrade with multiple upgrade groups, you can choose to immediately upgrade an upgrade group that has not been upgraded yet. This function is available only for scheduled upgrades.

To upgrade an upgrade group now:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Toggle Group List View**.
2. In the Group List view, click the Upgrade Group Now icon in the upgrade group row. Grid Manager immediately upgrades the selected group.

## Upgrading a Single Member Immediately

After the Grid Master has been upgraded, you can choose to immediately upgrade a specific member that has not been upgraded yet. This function is available only for scheduled Grid upgrades from NIOS 6.4.0 to a later release. You can upgrade a single member only when the Grid upgrade is paused, and you cannot upgrade the Grid Master, reporting appliance, and an offline member. Once the member has been manually upgraded, the appliance skips this member when its scheduled upgrade time is reached.

To upgrade a specific member now:

1. From the **Grid** tab, select the **Upgrade** tab.
2. Pause the upgrade.
3. Click **Toggle Member List View**, and select the member checkbox from the table.
4. From the Toolbar, click **Upgrade** -> **Upgrade Single Member**. Grid Manager immediately upgrades the selected member.

## Reverting a Single Member

During an upgrade from NIOS 6.4.0 to a later release, you can revert a specific member that has already been upgraded and is within its revert time window. The revert single member feature is useful when you want to troubleshoot issues, such as service outages, on a specific member after it has been upgraded. You can revert a member only when the Grid upgrade is paused, and you cannot revert the Grid Master, reporting appliance, and an offline member. If the upgrade is paused and you have reverted a member in an upgrade group that has already completed the upgrade, you must move the member to another upgrade group that has not been upgraded before you can proceed with the upgrade.

Once a member is upgraded, the appliance starts counting down and displays the time that is left for you to revert this member. You can revert the member before the revert time window expires. The default time window to revert a member is 24 hours. You can view the time that is left to revert the member in the Member List view, as described in Grid and Member Status below. You can also use the CLI commands [set default\\_revert\\_window](#) to configure the default revert time window for the Grid. For information about this command, refer to the *Infoblox CLI Guide*. Once a member exits the revert time window, you must revert the entire Grid in order to revert the member.



### Note

You may potentially lose some data when you revert a member. The appliance keeps information about DHCP leases and DNS records intact.

To revert a specific Grid member during a scheduled Grid upgrade:

1. From the **Grid** tab, select the **Upgrade** tab.
2. Pause the upgrade.
3. Click **Toggle Member List View**, and then select the member checkbox.
4. From the Toolbar, click **Revert** -> **Revert Single Member**.  
Grid Manager displays a message indicating that the revert process disrupts Grid services. Read the message carefully, and then click **Yes** to confirm your decision to revert the member. Be aware that when you revert a member, some changes made since the member was last upgraded may get lost.

## Upgrade Process

When an upgrade starts, Grid Manager checks if the nodes of an HA Grid Master have the same NIOS software version on their alternate partitions. If they do not have the same software version, the upgrade process stops. Grid Manager displays an error message and if it is a scheduled upgrade, Grid Manager deactivates the schedule as well. Otherwise, the upgrade process continues.



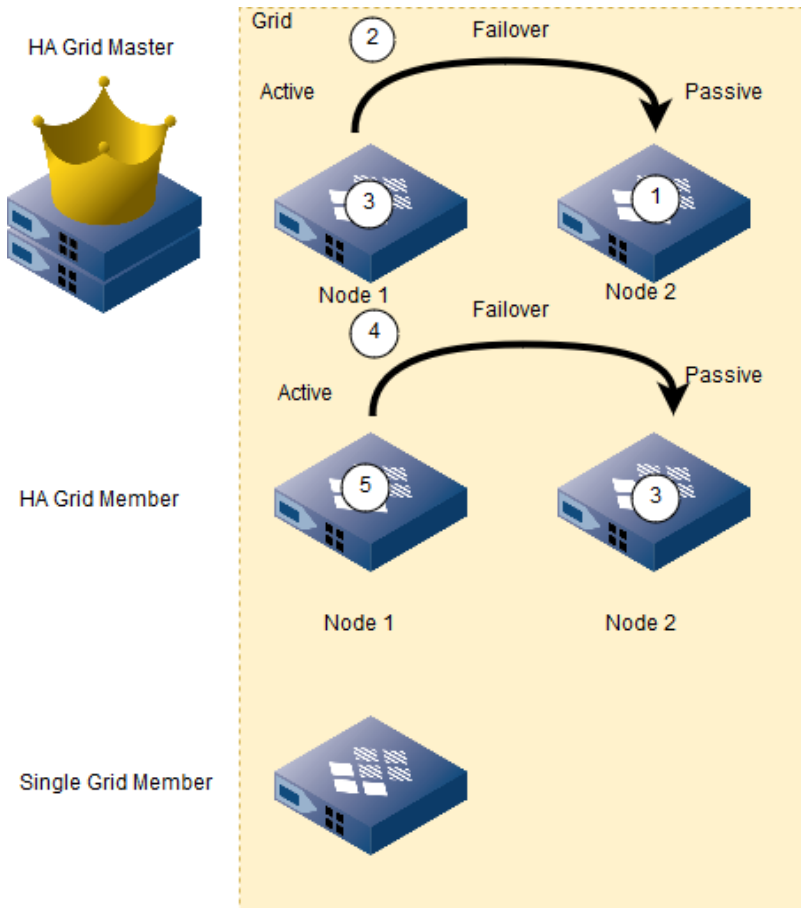
### Note

During the upgrade, you can view the status of the Grid Master in the serial console.

During the upgrade, if a Grid member has not completed its distribution, it automatically resynchronizes with the Grid Master after the Grid Master upgrade is complete.

Due to the nature of the upgrade sequence, HA pairs fail over during the upgrade. Therefore, be aware that the active and passive nodes reverse roles. The order in which Grid members upgrade, including when HA pairs fail over, is shown in the below figures Upgrade Sequence for an HA Grid Master and Grid Members and Upgrade Sequence for a Single Grid Master and Grid Members.

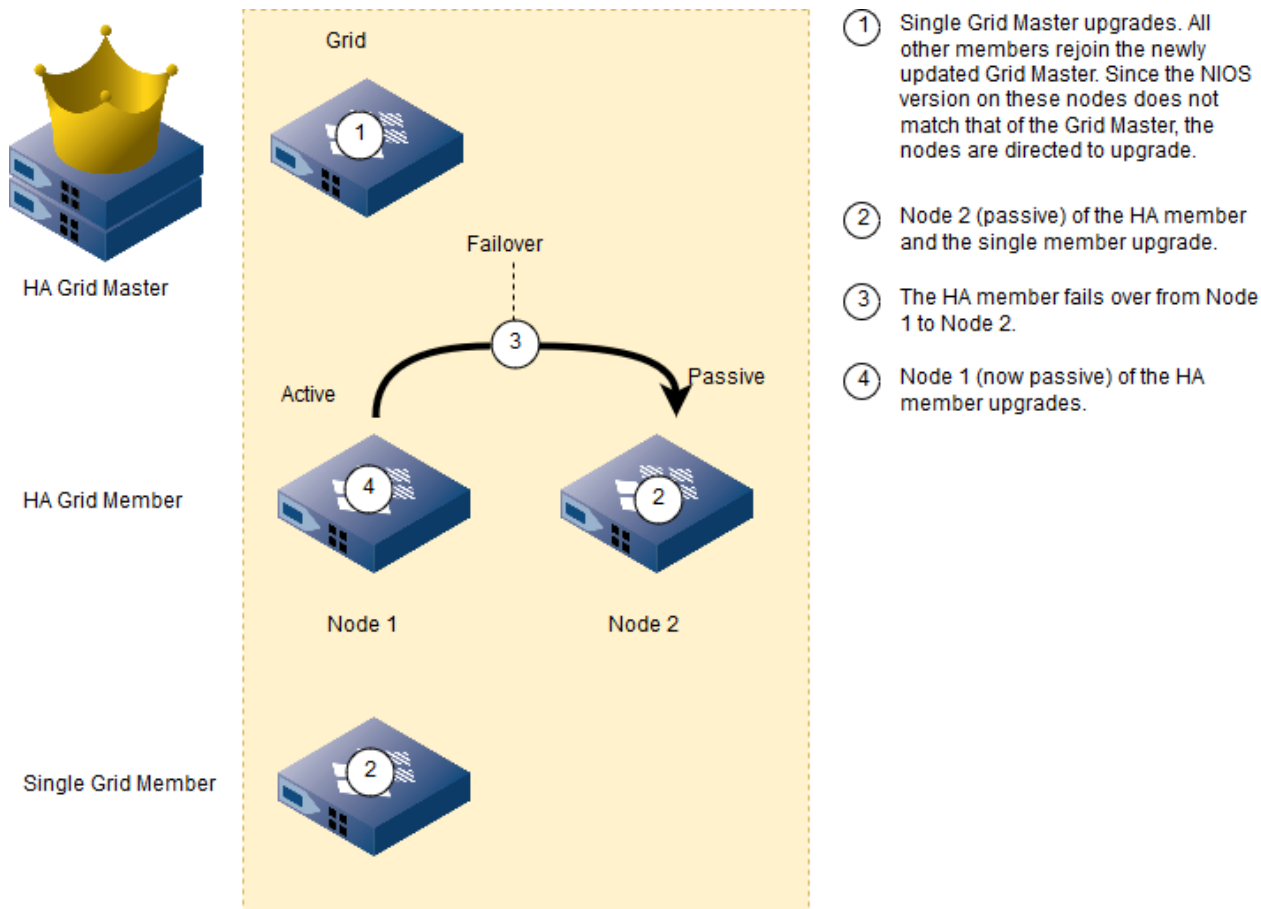
*Upgrade Sequence for an HA Grid Master and Grid Members*



- ① The passive node (Node 2) of the Grid Master upgrades.
- ② The Grid Master fails over from Node 1 to Node 2. At this point, the active Grid Master (Node 2) is using the upgraded code. All other nodes, including the passive node (Node 1) and all Grid members, rejoin the newly updated active node (Node 2). Since the NIOS version on these nodes does not match that of the active Grid Master, the nodes are directed to upgrade.
- ③ Node 1 (now passive) of the Grid Master upgrades. The passive node (Node 2) of the HA member and the single Grid member upgrade.
- ④ The HA Grid member fail overs from Node 1 to Node 2.
- ⑤ Node 1 (now passive) of the HA member upgrades.

**⚠ Note**  
Grid members that do not have the correct NIOS version on their alternate partitions due to an incomplete distribution automatically resynchronize the NIOS version with the Grid Master, and then upgrade.

*Upgrade Sequence for a Single Grid Master and Grid Members*



The Grid Manager session terminates when the HA Grid Master fails over from Node 1 to Node 2, or when the single Grid Master reboots and goes offline. During a scheduled upgrade, the Grid members that have not upgraded yet can join the Grid and function normally until their scheduled upgrade time. When the upgrade finishes, the upgrade schedule is set to inactive.

## Managing Upgrades

During an upgrade, Grid Manager displays a system message at the top of the screen indicating the Grid is being upgraded. After you start an upgrade, you can pause or resume it. For information, see [Pausing and Resuming Upgrades](#) and [Monitoring Distribution and Upgrade Status](#) below.

### Pausing and Resuming Upgrades

The following are some operational guidelines for performing an upgrade:

- You may not be able to perform certain administrative tasks during an upgrade.
- The Grid Manager session terminates when an HA Grid Master fails over from Node 1 to Node 2, or when a single Grid Master reboots and goes offline. You can log back in to the appliance after the upgrade.
- When you pause an upgrade, you can do the following
  - Change the sequence of the upgrade groups
  - Change the scheduled upgrade time for an upgrade group

To pause an upgrade, from the Grid Upgrade Status bar, click the Pause icon. When you pause an upgrade, Grid Manager displays a system message at the top of the screen indicating the upgrade is paused, until you resume the

upgrade. For information about the upgrade status of each member, see Monitoring Distribution and Upgrade Status below.

To resume an upgrade:

1. From the Grid Upgrade Status bar, click the Resume icon.
2. When the appliance displays a dialog box confirming that you want to resume the upgrade, click **Yes** to continue.

Members that have not completed or started upgrades that were scheduled at an earlier time resume the upgrade.

## Monitoring Distribution and Upgrade Status

During a distribution or an upgrade, Grid Manager displays the status of the distribution or upgrade in the status bar. It also displays the process status for each member. You can view the status in either the Member List view or Group List view from the **Grid** tab -> **Upgrade** tab.

When you perform a distribution or an upgrade, the status bar displays the overall Grid distribution status with a progress bar that describes the process being performed. The status bar also displays the number of members that have completed the distribution or upgrade.

A difference between a distribution and an upgrade process is that during an upgrade, the Grid Manager session terminates when an HA Grid Master fails over from Node 1 to Node 2, or when a single Grid Master reboots and goes offline. You can log back in to the appliance after the upgrade.

## Grid and Member Status

You can view the distribution and upgrade process status at the Grid and member level. To view the process status, from the **Grid** tab, select the **Upgrade** tab, and then click **Toggle Member List View**.

The status bar displays the status of the overall Grid process. It contains a progress bar that indicates the percentage of completion. It also shows the number of members that have completed the process.

Grid Manager displays the following information for each member:

- **Member:** The name of the Grid member.
- **Group:** The upgrade group to which the member belongs.
- **HA:** Indicates whether the member is an HA pair or not.
- **Status:** The current distribution or upgrade status. This can be Running (green) or Offline (red).
- **IPv4 Address:** The IPv4 address of the member.
- **IPv6 Address:** The IPv6 address of the member.
- **Running Version:** The NIOS software version that is currently running on the member.
- **Alternate Version:** Displays the NIOS software version to which the appliance can revert.
- **Distribution/Upgrade Status:** The current distribution or upgrade status. When the distribution or upgrade is in progress, Grid Manager displays a progress bar in this field to indicate the percentage of completion.
- **Hotfix:** The name of the hotfix that was last run on the member.
- **Status Time:** The date, time, and time zone of the status displayed.
- **Member Revert:** Indicates whether the member has been reverted or not. This appears only when the member has been upgraded from NIOS 6.4.0 to a later NIOS release.
- **Time to Revert:** The time (in HH:MM:SS format) left to revert a member. This appears only when the member has been upgraded from NIOS 6.4.0 to a later NIOS release.
- **Site:** The location to which the member belongs. This is one of the predefined extensible attributes. The appliance automatically refreshes the information in this panel.

## Upgrade Group Status

You can view the distribution or upgrade status of an upgrade group in the group list view. In this view, the distribution or upgrade status rolls up to the group level. You can expand an upgrade group to view the status of individual member. However, you cannot view detailed status of a selected member from this view.

To view the process status of an upgrade group, from the **Grid** tab, select the **Upgrade** tab, and then click **Toggle Group List View**. Grid Manager displays the following information for each member in an upgrade group:

- **Group:** The upgrade group to which the member belongs.

- **Member:** The name of the Grid member.
- **Status:** The current member status. This can be Running (green) or Offline (red).
- **IPv4 Address:** The IPv4 address of the member appliance.
- **IPv6 Address:** The IPv6 address of the member appliance.
- **Running Version:** The NIOS software version that is currently running on the member.
- **Distribution Status:** The current distribution status. For an upgrade group, Grid Manager displays a progress bar to indicate the overall percentage of completion. For a member, Grid Manager displays the state of the distribution process.
- **Timestamp:** The date, time, and time zone of the status displayed.

## Detailed Status

You can view detailed process information of a member during a distribution or an upgrade. To view detailed process information:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Toggle Member List View**.
2. Select a member and then click the Detailed Status icon.

Grid Manager displays a panel that shows the required steps during a distribution or an upgrade. It also displays a color indicator, next to each step, to indicate the current status of each step. The color indicator can be one of the following:

- Grey: The process has not started yet.
- Green: The process is complete.
- Blue: The distribution or upgrade that is in progress.
- Red: There is an error; Grid Manager displays a description of the problem.
- Yellow: A warning message.

When the selected member is an HA pair, Grid Manager displays the status information for both nodes. The panel remains open until you close it or select a different member.

## Lite Upgrades

A lite upgrade occurs when there are incremental changes to the software that do not require any upgrade to the database. The appliance can perform a lite upgrade only if the format of the database between the existing NIOS version and the upgrade version is the same.

In general, when you upgrade from a patch release to another patch release, you are performing a lite upgrade. In a lite upgrade, members can be running a different software version than the Grid Master. You can add objects, such as zones, networks, and resource records to the members that are running an older NIOS version. Replication of zones, networks, resource records, and DHCP leases is supported between the Grid Master and members. When you want to revert a member however, you must revert the entire Grid.

Whenever possible, the appliance uses the lite upgrade mode to speed up the upgrade process. You can always schedule a lite upgrade. Note that the appliance disables the testing function for lite upgrades because you do not need to test a lite upgrade for any database translation. For information about how to schedule an upgrade, see [Scheduling Upgrades in \*Upgrading NIOS Software\*](#).

## Viewing Software Versions

Before you upgrade, downgrade, or revert to a different NIOS software version, you can view the current software version that is running on the Grid, the NIOS image you have uploaded, and the available version to which you can revert. Grid Manager displays the software information in the **Upgrade** tab.

To view software information:

1. From the **Grid** tab, select the **Upgrade** tab.
2. Grid Manager displays the following in the Grid Version Information section:
  - **Running:** The NIOS software version that is currently running on the Grid.
  - **Uploaded:** The latest NIOS image file you have uploaded and is available for distribution.
  - **Distribution:** The NIOS software version used for distribution or is available for distribution.



- **Revert:** The NIOS software version to which the appliance can revert.
- **DistributionSchedule:** Displays the date and time of the next scheduled distribution.
- **UpgradeSchedule:** Displays the date and time of the next scheduled upgrade.

 **Note**

Grid Manager leaves a field empty when there is no available software for the specific function.

Grid Manager automatically refreshes the **Upgrade** tab with the latest information and displays the timestamp in the **Last Updated** field below the Grid Version Information section.

## Downgrading Software

Each Infoblox appliance model has a minimum required release of Infoblox software. Before downgrading an appliance, refer to the document, *Minimum Required Release Software for Hardware Platforms*, that shipped with your product. The downgrade procedure is for single independent appliances only. Infoblox does not support software downgrades for Grid members, but you can revert to the previous NIOS release (see the next section) on a Grid Master.

 **Note**

Although the downgrade process preserves license information and basic network settings, it does not preserve data. After you complete the downgrade procedure, all data in the database is lost.

To downgrade software on a single independent appliance running NIOS 4.0 or later:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Downgrade** from the Toolbar. Grid Manager displays a warning indicating that reverting to the current release is not possible once you start the downgrade. Read the warning carefully, and then click **Yes** to confirm your decision to downgrade.
2. In the *Choose file* dialog box, navigate to the downgrade image file, and then click **Open** to upload the file. The appliance uploads the file to the Grid Master. You cannot stop the downgrade process once you start it. Grid Manager displays the downgrade status in the status bar.

## Reverting the Grid to the Previously Running Software

You can revert the Grid to a version of software that was previously running on your NIOS appliance. The NIOS appliance stores the previous software version in its backup software partition. You can see if there is a software version to which you can revert and its version number in the Alternate Version column in the Grid Version Information section of the **Upgrade** tab. To view the software version, from the **Grid** tab, select the **Upgrade** tab. Note that once you start distributing a new NIOS version after an upgrade, you cannot revert to a previous NIOS version.

Be aware that when you revert to this software, changes made since the Grid was last upgraded are lost, including the new DHCP leases and other DNS changes.

To revert to a version of software previously running on a Grid or on an independent appliance or HA pair:

- From the **Grid** tab, select the **Upgrade** tab, and then click **Revert** -> **Revert Grid** from the Toolbar.

Grid Manager displays a warning indicating that the revert process disrupts Grid services. Read the warning carefully, and then click **Yes** to confirm your decision to revert.

## Applying Hotfixes

Infoblox periodically releases hotfixes that contain resolved issues. Only superusers can apply hotfixes through Grid Manager. When you install hotfixes through Grid Manager, you can apply them to the Grid Master and All Grid members, the Grid Master only, the Grid Master and Grid Master Candidates, or selected Grid members. This feature is supported on appliances running NIOS version 7.1 or later. Note that each hotfix addresses specific issues.

Infoblox recommends that you verify the hotfix before you apply it to ensure that it is the correct version.



After you apply a hotfix, Grid manager displays the hotfix status in the **Upload Status** bar. In addition, you can view the history of the most recent list of applied hotfixes in the *Hotfix History* dialog box.

 **Note**

A hotfix installation may fail if there is a mismatch in the NIOS software versions or if a hotfix image fails to meet the software or hardware restrictions.

To apply a hotfix:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **ApplyHotfix** from the Toolbar and select one of the following:
  - **ToGridMasterandallGridMembers**: Click this to apply hotfix to the Grid Master and all Grid members.
  - **ToGridMaster**: Click this to apply hotfix to the Grid Master only.
  - **ToGridMasterandGridMasterCandidates**: Click this to apply hotfix only to the Grid Master and Grid Master candidates.
  - **ToSelectedGridMembers**: Click this to apply hotfix to the selected Grid member. This option is available only after you have selected a Grid member or members.
2. In the *ApplyHotfix* dialog box, click **Select** and navigate to the hotfix image file you want to upload. Click **Open** to select the file, and click **Upload**.

 **Note**

If you have already installed a hotfix and subsequently try to upgrade or downgrade NIOS, the appliance displays a warning message because resolved issues in the hotfix will no longer be valid if the upgrade or downgrade version does not contain the original hotfix.

The appliance applies the hotfix to the targeted appliance(s). You cannot stop the hotfix process once you start it. Grid Manager displays the hotfix status in the **Upload Status** bar.

## Viewing Hotfix History

You can view hotfix history on selected members. To view the hotfix history:

1. From the **Grid** tab, select the **Upgrade** tab, click the Show Hotfix History icon in the **Hotfix** column.
2. The *Hotfix History* dialog box displays the following information:
  - **Date**: The timestamp when the hotfix was applied.
  - **Hotfix**: The name of the hotfix that was last run on the member. The name of the hotfix image file that was applied to the member.
3. Click **Close** to close the *Hotfix History* dialog box.

## Using Database Snapshots

Infoblox recommends that you regularly create database snapshots that helps in mitigating the impact of any user errors in the NIOS configuration. Whenever there is an error in the NIOS configuration, you can roll back the database (onedb) to the snapshot that you have created earlier. This is potentially faster than restoring the database using the backup file. The following sections describe how to create snapshots and roll back snapshots:

- Creating Database Snapshots
- Rolling Back Database Snapshots

## Creating Database Snapshots

You can create a database snapshot periodically and save it locally on the appliance. You can create only one copy of the database snapshot at any given time. Each time you create a new snapshot, it overrides the previous snapshot. When you create a database snapshot, only data that is currently saved in the database is included in the snapshot. Only


superusers are allowed to create a database snapshot. Creating database snapshots does not have any impact on the scheduled local backup. The snapshot saved locally on the appliance will not be affected if an HA failover occurs, but the snapshot is deleted during a Grid Master Candidate promotion, upgrade, downgrade, or if you restore the appliance using the backup file. Note that creating a database snapshot fails when any of the following operations are in progress:

- Rolling back database snapshot.
- When another database snapshot is being created.
- Restoring backup files. To create a snapshot:

1. From the **Grid** tab -> **Grid Manager** tab, click **Snapshot** -> **Create** from the Toolbar.
2. In the *Create Snapshot* dialog box, enter information about the snapshot in the **Comment** field and click **OK**.
3. In the *Create Snapshot* confirmation dialog box, click **Yes**.

## Rolling Back Database Snapshots

If you have already created a snapshot, you can roll back the database to the snapshot to recover from any user errors in the NIOS configuration. Rolling back the database (onedb) to the snapshot is potentially faster than restoring the database from a backup file. Rolling back to the snapshot changes only the database, not any other transient data, such as file distribution, reporting splunk, discovery data, etc. Note that you might lose the data that you have saved after creating the snapshot. Only superusers can perform the rollback operation.

 When reporting cluster is enabled, rolling the database back to the snapshot that was created when reporting was in the single-indexer mode might cause some data loss. The database settings will reset reporting to the single-indexer mode, and the indexer might not have all the data indexed since reporting cluster was enabled.

To roll back the database to the snapshot:

1. From the **Grid** tab -> **Grid Manager** tab, click **Snapshot** -> **Rollback** from the Toolbar.
2. The *Rollback Snapshot* dialog box displays the following information:
  - **Snapshot Time:** The date and timestamp when the snapshot was created.
  - **Comment:** Information about the snapshot.
3. Click **OK**.

## Downloading Support Bundles

When you need assistance troubleshooting a NIOS appliance, you can log in to the appliance as a superuser, download the support bundle of the appliance, and then send it to Infoblox Technical Support for analysis. A support bundle is a tar.gz file that contains configuration files and the appliance system files. You can download a support bundle for an independent appliance and for each member in a Grid. When you download a support bundle for an HA pair, it includes the files of both nodes in the HA pair.

By default, the appliance includes the following files in the support bundle: core files, current logs, and rotated logs. Because core files can be quite large and take a significant amount of time to download, Infoblox recommends that you include core files in the support bundle only when requested by Infoblox Technical Support. You can include all the core files or specific core files in the support bundle when Infoblox Support requests them. Note that the option of downloading only specific core files is supported on appliances running NIOS version 8.0 or later. If your appliance supports multiple primary servers for DNS authoritative zones, you can select to include cached zone data in the support bundle. This data is not included in the bundle by default.

You can override the default support bundle download time which is 1200 seconds. To do so, go to **Grid** -> **Grid Manager** -> **Grid Properties** -> **Edit** -> **Advanced**, and set a custom value for **Support Bundle Download Timeout**. You can also use the `set support_timeout` command from the Administrative Shell.

To download a support bundle:

1. From the **Grid** tab, select a *member* checkbox, and then click **Download** -> **Support Bundle** from the Toolbar.
2. In the *Download Support Bundle* dialog box, select the files you want to include in the support bundle, and then click **OK**.

- **Core Files:** Infoblox recommends that you include these files only when requested by Infoblox Technical Support.
  - **Current Logs:** Infoblox recommends that you always include these files in the support bundle.
  - **Rotated Logs:** These are rotated logs that contain historical information.
  - **Cached DNS Recursive Data:** Select this to include cached DNS recursive data in the support bundle. This is deselected by default.
  - **Cached Zone Data:** Select this to include cached DNS zone data in the support bundle. This is deselected by default.
  - **Subscriber Services Data:** If you have configured Infoblox Subscriber Services, you can select this to include subscriber services data in the support bundle. This is deselected by default. For information about Infoblox Subscriber Services, see [Infoblox Subscriber Services](#).
  - **Cloud Discovery Data:** Select this to include cloud discovery data in the support bundle. This is deselected by default.
  - **Files from the previously installed NIOS version:** Select this to include files from the previously installed NIOS version in the support bundle. This is deselected by default.
  - **Discovery SNMP Logs:** Event logs related to device discovery SNMP probes of routers, switches and other network infrastructure devices.
  - **Core Files:** Infoblox recommends that you include these files only when requested by Infoblox Technical Support.
    - **All Core Files:** Select this to include all core files in the support bundle.
    - **Select Core Files:** Select this to include only specific core files in the support bundle. Click the Add icon in the table. Grid Manager displays the *Core Files Selector* dialog box from which you can select specific core files. You can delete a core file by selecting its checkbox and clicking the Delete icon. Note that on an HA pair, you can select the core files of the active node only.
3. Navigate to the location you want to save the file and change the file name. Do not change the .tar.gz file extension in the file name.
  4. Send this file to Infoblox Technical Support.

## Guidelines for Upgrading, Backing Up, and Restoring Data

You should take into consideration the impact on scheduled and approval tasks when you perform any of the following:

- When you upgrade from previous releases to NIOS 6.7 and later, the appliance converts all valid punycode data to IDNs for DNS resource records and DNS zones. When you have a fresh installation of NIOS 6.7 and later, the appliance converts all valid punycode data to IDNs for DNS zones only. It retains punycode data for resource records.
- Upgrade the NIOS software: In a full upgrade, all scheduled and approval tasks are deleted. In a lite upgrade, scheduled and approval tasks are not deleted.
- Back up the NIOS database: All scheduled and approval tasks are backed up for troubleshooting purpose.
- Restore the database: Scheduled and approval tasks are not restored.
- Promote a Grid member to a Grid Master: After the promotion, all scheduled and approval tasks that are past due are executed immediately.
- Revert the NIOS software image: After the revert, all scheduled and approval tasks that are past due are executed immediately.
- Restore data from the Recycle Bin: To restore a deleted parent object (such as a network) that contains a child object (such as a DHCP range) associated with a scheduled or approval task, you must first delete the scheduled or approval task for the child object.

### Related topic

[Scheduling Tasks](#)

# Guidelines for Upgrading the Reporting and Analytics Solution

When you upgrade from a previous NIOS release to NIOS 7.3.x and later releases, you will notice that some of the reporting features and terminologies have changed. For example, searches and reports in previous NIOS releases are now reports and dashboards respectively in the new reporting solution. In addition, your custom reports might be affected. Infoblox recommends that you take some time to explore the new user interface and get familiar with the terminologies.

You can continue to use the Infoblox predefined reports from previous releases in the new interface and customize them to meet your specific requirements, or you can create new custom reports from the ground up using a powerful search pattern.



## Note

Infoblox Reporting and Analytics integrates with Splunk to deliver an enhanced reporting interface so you can create dashboards, reports, and alerts. This chapter attempts to explain all reporting functionality you can perform through the enhanced interface. However, you may need to refer to the Splunk documentation for certain functionality as indicated in specific sections of this chapter. In addition, some functions and capabilities referenced in the Splunk documentation, such as setting up custom Python scripts, are not available or applicable to the Infoblox Reporting and Analytics as some Splunk functionality in the Infoblox product may be limited or modified by Infoblox. Infoblox does not represent or warrant that Infoblox Reporting and Analytics will function in accordance with the Splunk documentation. Infoblox is also not responsible for the accuracy of the Splunk documentation. For Infoblox Reporting and Analytics technical support, contact Infoblox Technical Support. DO NOT contact Splunk.

When you upgrade from a previous NIOS release to NIOS 7.3.x and later, there are some significant changes to the Reporting and Analytics solution. Some of the important changes are as follows:

- **Terminology:** The following table lists the terminology differences when you upgrade:

### Reporting Terminology Changes

Pre-NIOS 7.3.0 Release	NIOS 7.3.0 and later
Searches	Reports
Reports	Dashboards

- **Object Management:** NIOS no longer manages reporting objects such as searches, smart folders, alerts, and reports. You will not be able to perform operations such as global search, quick filtering, bookmarking, and others for these objects. You can now manage these objects through the new user interface. Smart folders are migrated after an upgrade. However, data in the smart folders is not migrated, and all filters for the smart folders are reset to default.
- **Permissions:** Permissions for all reporting objects are migrated to the new Reporting and Analytics solution and managed through the new user interface after an upgrade. You may see the new built-in role, **Everyone**, when configuring Reporting permissions. For best practices, do not alter permissions for this new built-in role. Note that the **Reporting Dashboard** and **Reporting Search** global permissions have been removed. If an admin group or admin role was granted these permissions before an upgrade, the permissions will still be displayed after an upgrade. However, they won't take any effect. The **Grid Reporting Properties** permission is retained. In addition, reporting object permissions for dashboards and searches (including global dashboards and searches) are migrated. These object permissions are retained for applicable migrated users. If permissions were granted to a specific admin group for a dashboard or search before an upgrade, only these admin users and superusers have permissions to access the migrated dashboard and report after an upgrade. If a limited-access user group is created through the new interface after the upgrade, users in this admin group will not be able to access the dashboard and report even if they are granted access to the **Infoblox Reporting and Analytics App**. Superusers

must explicitly grant permissions to this limited-access admin group for users in this group to access the dashboard and report. For more information, see [Administrative Permissions](#).

- **Navigation and Visualization:** Navigations for some reporting functions, such as searches, alerts, email and page settings, and email PDF delivery, have changed. You can navigate through the new user interface to get familiar with the changes in this release. In addition, all predefined reports might look different than the traditional ones depending on your filtering configuration. The *Grid Reporting Properties* editor and **Groups** tab are moved under the **Administration** tab → **Reporting** tab.  
Note that the Bookmarked groups are migrated after an upgrade. The bookmarked group navigates to the **Administration** tab → **Reporting** tab → **Groups** tab.
- **Extensible Attributes:** The reports that supported filtering and grouping by multiple extensible attributes are migrated to the new interface with filtering and grouping only by the extensible attribute **Site**. You must clone the dashboard, add filter inputs and modify the view XML to support additional extensible attributes. For information about Editing the XML Source Code of a Dashboard, see [About Dashboards](#).
- **Searches and Reports:** Only NIOS system and global reports and searches are migrated to NIOS 7.3.0 and later versions as dashboards and reports respectively. All user private reports and searches are dropped. In addition, bookmarked reports and searches are not migrated to 7.3.x release. If you want to keep any customization for the user private dashboards and reports, do one of the following:
  - Create global dashboards and reports using the same settings.
  - After an upgrade, you can clone the corresponding migrated system or global dashboards and reports, and then reconfigure the original settings, such as filters and scheduling in the new user interface.
- **Custom Search:** You can create your own search pattern and save it as a dashboard or report. For information about About Searches, see [Home Dashboards](#).

After completing the NIOS upgrade successfully, you configure the *Grid Reporting* properties and remote server (FTP, SCP, or TFTP) to export search results. For information, see [Grid Reporting Properties](#) and [Configuring an External Server for Search Result Exports](#).

## Related topic

[Infoblox Reporting and Analytics](#)

## Guidelines for Scheduling Full Upgrades

When scheduling a full upgrade in NIOS, the Grid Master replicates the following features to the Grid members, including to those members that have not been upgraded:

- DNS resource records
- DNS zones
- DNS views
- Name server groups
- Shared record groups
- IPv4 and IPv6 host addresses
- Roaming hosts
- IPv4 and IPv6 networks
- IPv4 and IPv6 shared networks
- Fixed addresses
- DHCP ranges
- DHCP failover association
- DHCP option spaces
- DHCP options
- DHCP filters
- MAC filter items
- Blacklist & NXDOMAIN rules
- DNSSEC key pairs
- DNSSEC import keyset operation
- Signed and unsigned zones

- DNSSEC rollover KSK and ZSK operations.

You can perform the following tasks during an upgrade:

- Upgrade a specific member during the scheduled Grid upgrade. For information about how to upgrade a single member during a scheduled Grid upgrade, see [Upgrading a Single Member Immediately in \*Upgrading NIOS Software\*](#).
- Revert a single member that has already been upgraded to troubleshoot issues, such as service outages, on that specific member. The upgrade of that member can then be rescheduled. For more information, see [Reverting a Single Member in \*Upgrading NIOS Software\*](#).
- Clear authentication cache and authentication records.
- Perform AD (Active Directory) configurations. Note that the keytab file must be uploaded before the upgrade starts.

Note the below restrictions when scheduling a full upgrade:

- Any action that requires a service restart for configuration changes to take effect are not recommended and can result in upgrade issues.
- Do not add, modify, or delete an NS group.
- Do not add, modify, or delete manually created NS records.
- Do not add, modify, or delete a zone.
- Do not assign or unassign an NS group to a zone.
- Do not change the NS group assigned to a zone.
- Do not change the host name of the Grid members that are assigned to a zone if the members have not been upgraded, have been reverted, or are in the revert time window.
- Do not restart DNS and DHCP services or schedule a restart for these services on Grid members that have not been upgraded. For more information about Restarting Groups, see [Restarting Services](#).
- Do not add, delete, or modify a DHCP range, a filter, or a fixed address.
- Do not modify the settings for automated mitigation of phantom domain attacks using the CLI commands on a Grid member until the member has completed the upgrade.

The Grid Master and Upgrade Groups can be scheduled to upgrade at different times in order to limit service impact. However, upgrades need to be performed within a limited window of time (i.e. within a couple of days). If an upgrade spans nine or more days, a warning is displayed in the NIOS UI.

NIOS does contain checks and rules to ensure data integrity that can cause undesirable results during the upgrade process. However, when scheduling a full upgrade, the following rules and behavior has to be noted and followed to ensure a seamless upgrade:

- Do not modify member properties for the following: DNS, DHCP, TFTP/HTTP/FTP, bloxTools, Captive Portal, Reporting, and load balancing until the member has completed the upgrade and exited its revert time window.
- Do not delete DNS views until the entire Grid upgrade is complete.
- Do not delete DNS zones and IPv4 and IPv6 networks that are under Microsoft Management until the managing member of the Microsoft servers has completed its upgrade and exited its revert time window.
- Synchronization between load balancers and the appliance is disabled until the load balancer managing member has completed its upgrade. Do not change the managing member during the upgrade.
- Do not add, modify, or delete network views, rulesets, and DNS64 synthesis groups until the entire Grid upgrade is complete.
- Replication of Grid and member DNS and DHCP properties is not supported.
- Do not create additional named Access Control Lists (ACLs) until after the entire Grid has been upgraded. For information about named ACLs, see [Configuring Access Control](#).

During a scheduled full upgrade, the Grid Master skips Grid members that do not complete their NIOS upgrade within 10 minutes, the default upgrade policy time, and moves to the next Grid member within the upgrade schedule.

During a scheduled full upgrade, do not perform the following tasks on a Grid member that has not been upgraded yet:

- Import the DHCP lease history file
- Use the DHCP expert mode configuration feature
- Clear the NAC authentication cache of a DHCP member
- Set the time zone for a Grid member
- View the capacity report of a Grid member

- Test the email configuration settings of a Grid member
- Check whether an IPv6 address is already configured on a Grid member

When scheduling a full upgrade from a previous NIOS release to a release that includes the DHCP fingerprint detection feature, the following rules apply until the entire Grid has been upgraded:

- DHCP fingerprint detection is disabled
- Do not add DHCP fingerprint filters
- Do not apply DHCP fingerprint filters to any DHCP address range

When scheduling a full upgrade from a previous NIOS release to a release that includes the multi-primary zone feature, the following rules apply until the entire Grid has been upgraded:

- Do not configure multiple primary servers for an authoritative zone or configure a name server group that contains multiple primary servers.
- Do not assign or unassign a Grid member to an authoritative zone or name server group.
- Do not change the stealth state of an authoritative zone or name server group.

When scheduling a full upgrade from a previous NIOS release to a release that includes the Infoblox Threat Protection feature, do not perform the following on a Grid member until the member has completed the upgrade:

- Start or stop the Threat Protection and DNS services.
- Activate a ruleset.
- Perform any threat protection related tasks such as adding custom rules and activating rulesets.

Before scheduling a full upgrade from a previous NIOS release to a release that includes the IPv6 Grid feature, the following rules apply:

- If the Grid has an HA Master or HA member and if it is configured with IPv6 VIP address, IPv6 addresses must be configured for both node 1 and node 2.
- Both the Grid Master and the Grid Master Candidate must have the same type of network connectivity.
- The current configuration and database must be backed up.
- If the subscriber site has HA and the HA passive node is the first to upgrade, the data repository connectivity uses the IPv4 protocol for the site members. If you want the data repository to be connected over the IPv6 protocol, you must stop and restart the subscriber service in the upgraded Grid. The subscriber data is lost when the service is stopped and restarted. It is recommended to stop/start the service of each member at a time to synchronize the subscriber cache with the next active member on the same site.

When scheduling a full upgrade from a previous NIOS release to a release that includes the Secure Dynamic Updates feature, the following rules apply until the Grid has completed the upgrade:

- All dynamic updated records are labelled as static records. Infoblox suggests to enable this feature only after all records are changed to Dynamic.
- NIOS tags the RRsets that are not auto-generated as static records. For information about Secure Dynamic Updates, see [Secure Dynamic Updates](#).

When scheduling a full upgrade that includes the DNS Traffic Control feature, the following rules apply until the entire Grid has been upgraded:

- Do not add an SNMP health monitor.
- Do not configure the All available load balancing method for a DTC pool.
- The record types are reset to default record types (A and AAAA records) and do not modify the record types for an LBDN.

## Upgrading Parental Control at DNS Cache Acceleration

Upgrading Infoblox subscriber services parental control at DNS Cache Acceleration using cached domain and subscriber data has the following restrictions:

- Upgrade subscriber services using a staged upgrade. This does not affect subscriber data.
- You must update parental control category data download credentials after the upgrade.
- When you upgrade, designate a few members per site to run garbage collection as subscriber services does not perform garbage collection.



- Restrictions when upgrading subscriber sites:
  - You cannot add or remove members from a site during an upgrade.
  - You cannot stop or start a subscriber secure service during an upgrade.
  - You cannot change any subscriber service configuration during an upgrade.

## Microsoft Management Rules

On a member that synchronizes data with Microsoft DNS and DHCP servers, the following functions are deactivated during an upgrade:

- Synchronization of Microsoft DNS and DHCP data
- Rotation of Microsoft logs
- Start and stop of Microsoft servers
- Releases of DHCP leases from a Microsoft DHCP server

Deactivation of these functions does not affect data on the Microsoft servers. After the upgrade, the member automatically restarts the synchronization of Microsoft data.

On a member that synchronizes data with Microsoft DNS and DHCP servers, the following rules apply:

- Do not modify the managing member if the old and new members have not been upgraded and have not exited their revert time windows.
- Do not add, modify, or delete zones, IPv4 DHCP ranges, and IPv4 networks until the managing member has been upgraded and exits the revert time window.
- Do not add, modify, or delete DNS resource records if the associated zone is managed by a Microsoft server and the managing member is still in its revert time window.
- Do not add, modify, or delete fixed addresses that are assigned to a Microsoft server and the managing member is still in its revert time window.
- Wait until the new managing member is upgraded to configure it as a DNS primary or secondary.

## DHCP Expert Mode Upgrade

Enabling DHCP expert mode allows administrators to directly manipulate sections of the DHCP configuration file. In this mode, all built-in protections and error checking normally provided by Infoblox are bypassed.

Because these protections are removed, Infoblox is unable to provide support for DHCP while in the DHCP expert mode except to confirm that administrator changes to the configuration file were written. The integrity of the configuration file when the DHCP expert mode is enabled is entirely the responsibility of the administrator. If Infoblox support for DHCP is required, first disable the DHCP expert mode and then reproduce the issue.

Infoblox strongly discourages the use of DHCP expert mode. Consider using it only after discussing the situation with Infoblox Support.

## Backing Up and Restoring Configuration Files

Infoblox recommends that you regularly back up your configuration files and/or discovery database files. You can back up your system files locally on the appliance or to your management system, or use TFTP (Trivial File Transfer Protocol), FTP (File Transfer Protocol), or SCP (Secure Copy) to back them up to a remote server. Backing up and restoring the configuration files using TFTP, FTP, or SCP is supported on both IPv4 and IPv6 communication protocols. You can select to back up files manually or schedule automatic backups for a later date.

To avoid missing a backup when a remote server is unavailable during a scheduled automatic backup, you can choose to save files locally on your appliance while backing up to the remote server. Both the local and remote backup files share the same date because NIOS saves these files from the same backup. The backup file is a .tar.gz file that contains the configuration settings, data set, and TFTP files. Note that the local backup contains only the Grid backup. It does not contain backups for reporting or NetMRI.





## Note

While the content of the backup file in plaintext it does not contain any plaintext representation of any passwords. These are either encrypted with AES-256 or salted hashed with SHA-128.

You may also schedule automatic backups of the discovery database, which consists of the complete discovery data for networks and network devices such as core, distribution and edge routers, enterprise switches, security devices, and end host devices. NIOS backs up the discovery database in a .tar.gz file, with the raw discovery data formatted as an XML file.



## Note

Infoblox recommends that you backup the configuration after you convert a Grid to a different mode. Restoring the old backup by performing a forced restore, may prevent the Grid members from rejoining the Grid Master after the restore.

The following sections describe how to use the backup and restore functions:

- Backing Up Files
- Automatically Backing Up Data Files
- Manually Backing Up Data Files
- Restoring Backup Files
- Downloading Backup Files from a Different Appliance



## Note

Infoblox highly recommends that you always back up the current configuration file before upgrading, restoring, or reverting the software on the appliance. If you are performing these operations on appliances licensed for Discovery and that perform discovery, the discovery database can be backed up and restored using the same mechanisms.

## Backing Up Files

You can back up system files and discovery databases periodically and on demand. You can then restore the files on the same appliance or on a different appliance. For information about restoring files, see Restoring Backup Files below. You can configure the appliance to automatically back up the files on a weekly, daily, or hourly basis.

Infoblox recommends that you back up the system files during off-hours to minimize the impact on network services. By default, the automatic backup function is turned off. You must log in with a superuser account to back up files.

You can back up system configuration and/or discovery database files to the following:

- A local directory
- The management system that you use to operate the appliance
- A TFTP server
- An FTP server. This option requires that you have a valid username and password on the server prior to backing up files.
- An SSH server that supports SCP. This option requires that you have a valid username and password on the server prior to backing up files.

## Local Backup

You can store a backup file on the appliance itself. However, Infoblox recommends that you store backup files in an alternate location. When you back up the system files locally, the appliance uses the following format to name the file: `BACKUP_YYYY_MM_DD_MM.tar.gz`. For example, a file name of `BACKUP_2013_11_30_23_00` means that the file is backed up on November 30th, 2013 at 11:00 PM.

The appliance can save up to 20 configuration files, regardless of how often the files are saved (weekly, hourly, or daily). Ensure that you take the size of the configuration file into consideration when backing up files because the storage limit on an appliance is 5 Gb (gigabytes). If your configuration file is 500 Mb (megabytes), then the appliance can store 10 configuration files. When uploading configuration files on to a TFTP, FTP, or SCP server, you must consider the file size on that server as well.

## Using TFTP

TFTP is a client-server protocol that uses UDP as its transport protocol. It does not provide authentication or encryption, therefore it does not require a username or password.

When you back up the system files to a TFTP server, you select the backup file you want to download, enter the name in which the file is stored on the TFTP server and the server IP address.

## Using FTP

FTP is a client-server protocol used to exchange files over TCP-based networks. The appliance, as the FTP client, connects to a remote FTP server that you identify. When you use FTP to back up the system files, the password and file contents are transmitted in clear text and may be intercepted by other users.

When you back up the system files to an FTP server, the appliance, as the FTP client, logs on to the FTP server. You must specify the username and password the appliance uses to log on to the FTP server. The user account must have write permission to the directory to which the appliance uploads the backup file.

## Using SCP

SCP is more secure than TFTP and FTP. It uses the SSH protocol to provide authentication and security. You can use SCP to back up the NIOS system files to a server running SSHv2.

When you use SCP to back up the system files to an SSH server, you must specify the username and password the appliance uses to log on to the server. Note that you must use either "password" or "Password" in the SCP password prompt because the appliance does not recognize "PASSWORD" in the prompt. Therefore, ensure that you customize the SCP password prompt to say "Enter your password" or "Enter your Password." Otherwise, the SCP backup will fail. The user account must have write permission to the directory to which the appliance uploads the backup file. In addition, make sure that you enter the correct IP address of the SSH server; the appliance does not check the credentials of the SSH server to which it connects.



### Note

The SCP protocol uses SSH for data transfer and thus provides the same authentication and security as SSH. SCP uses LAN1 regardless of whether the MGMT port is enabled or not.

## Automatically Backing Up Data Files

Infoblox recommends that you regularly back up your configuration files and/or discovery database files. The easiest way to accomplish this task is to configure the appliance for scheduled automatic backups of the NIOS configuration files.

When you automatically back up a configuration file on the appliance, the file is named in the format

<GRIDNAME>\_YYYY\_MM\_DD\_HH.MM.tar.gz. The default time for an automatic backup is 3:00 AM. Infoblox recommends scheduling configuration file backups to take place during the slowest period of network activity. You can choose a schedule for when and how often files are backed up: weekly, daily, or hourly.

If a Grid has a discovery member, you may also schedule automatic backups of the Discovery database, which consists of the complete discovery data for networks and network devices such as core, distribution and edge routers, enterprise switches, security devices, and end host devices. NIOS backs up the Discovery database in a .tar.gz file, with the raw Discovery data formatted as an XML file. For information on discovery features and requirements, see the chapter [Infoblox Network Insight](#).

To automatically back up a database file on an independent appliance or Grid Master:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Backup** -> **Schedule Backup** from the Toolbar.
2. In the *Schedule Backup* dialog box, select the destination of the backup file from the **Backup to** drop-down list:
  - **TFTP**: Back up system files to a TFTP server.
  - **Keep local copy**: Select this to also save a local copy of the backup file on your appliance. This is disabled by default. The local backup contains only the Grid backup, it does not contain backups for reporting and NetMRI. Note that when you select this, the total backup time will increase.
    - **IP Address of TFTP Server**: Enter the IP address of the TFTP server to which you want to back up the system files.
    - **Directory Path**: Enter the directory path of the file. For example, you can enter `/archive/backups`. The directory path cannot contain spaces and backslash (\). The folder or directory you enter here must already exist on the specified server. Do not include the file name in the directory path.
    - **Recurrence**: Select how often you want to back up the files. You can select **Weekly**, **Daily**, or **Hourly** from the drop-down list. When you select **Weekly**, complete the following:
      - **Every**: Choose a day of the week from the drop-down list.
      - **Time**: Enter a time in the hh:mm:ss AM/PM format. You can also click the clock icon and select a time from the drop-down list. The Grid Master creates a backup file on the selected day and time every week.  
When you select **Daily**, enter a time in the hh:mm:ss AM/PM format. You can also select a time from the drop-down list.  
When you select **Hourly**, complete the following:
        - **Minutes after the Hour**: Enter the minute after the hour when the Grid Master creates a backup file. For example, enter 5 if you want the Grid Master to create a backup file five minutes after the hour every hour.
    - **Disable Scheduled Backup**: Select this if you want to disable automatic backups from occurring now.  
You can still save the settings for future use.
  - **FTP**: Back up system files to an FTP server.
  - **Keep local copy**: Select this to also save a local copy of the backup file on your appliance. This is disabled by default. The local backup contains only the Grid backup, it does not contain backups for reporting and NetMRI. Note that when you select this, the total backup time will increase.
    - **IP Address of FTP Server**: The IP address of the FTP server.
    - **Directory Path**: Enter the directory path of the file. For example, you can enter `/archive/backups`. The directory path cannot contain spaces and backslash (\). The folder or directory you enter here must already exist on the specified server. Do not include the file name in the directory path.
    - **Username**: Enter the username of your FTP account.  
Note that if you have configured AD server for authentication, you must specify "domain name\username".
    - **Password**: Enter the password of your FTP account.
    - **Recurrence**: Select how often the scheduled backups should occur. You can select **Weekly**, **Daily**, or **Hourly**. For information, see TFTP.
    - **Disable Scheduled Backup**: Select this if you want to disable automatic backups from occurring now, but want to save the settings for future use
  - **SCP**: Back up system files to an SSH server that supports SCP.
  - **Keep local copy**: Select this to also save a local copy of the backup file on your appliance. This is disabled by default. The local backup contains only the Grid backup, it does not contain backups for reporting and NetMRI. Note that when you select this, the total backup time will increase.
    - **IP Address of SCP Server**: The IP address of the SCP server.

- **Directory Path:** Enter the directory path of the file. For example, you can enter **/archive/backups**. The directory path cannot contain spaces and backslash (\). The folder or directory you enter here must already exist on the specified server. Do not include the file name in the directory path.
  - **Username:** Enter the username of your SCP account.  
Note that if you have configured AD server for authentication, you must specify "domain name\username".
  - **Password:** Enter the password of your SCP account.
  - **Recurrence:** Select how often the scheduled backups should occur. You can select **Weekly**, **Daily**, or **Hourly**. For information, see the TFTP section.
  - **Disable Scheduled Backup:** Select this if you want to disable automatic backups from occurring now.  
You can still save the settings for future use.  
Note that when you select **FTP** or **SCP**, ensure that you have a valid user name and password on the server prior to backing up the files.  
New status types such as **Upload keys triggered**, **Upload keys in progress**, **Upload keys done** are displayed in the *Reporting Restore* dialog box also.
  - **Grid Master (Local):** Back up to a local directory on the Grid Master. This is the default.  
By default, the Grid Master generates a backup file and saves it locally in its own storage at 3:00 AM daily. Be aware that backing up the Grid and saving it locally on an hourly basis increases the turnover of files stored on the Grid Master. Backing it up hourly to a remote server increases the overall amount of traffic on your network.
3. If the Grid has a discovery member, Grid Manager displays the **NIOS data** and **Discovery data** checkboxes. You can select the **NIOS data** checkbox, to back up NIOS configuration data for the Grid and select the **Discovery data** checkbox, to back up discovery data for the Grid.  
If the Grid has a reporting member, Grid Manager displays the **Infoblox Splunk App** checkbox. You can select the **Infoblox Splunk App** checkbox, to back up Splunk application reporting data.
  4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Manually Backing Up Data Files

You can manually back up a NIOS data file in addition to scheduling your backups. If a Grid has a discovery member, you can also manually back up the current discovery database. Doing so backs up the complete discovery database that is resident on the Consolidator appliance, which is a member of the Grid. Keep in mind that discovery processes may be taking place on the associated NIOS appliances licensed for that task. NIOS will temporarily suspend the Discovery service while the backup is being retrieved from the Consolidator appliance.

To back up manually:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Backup -> Manual Backup** from the Toolbar.
2. In the *Backup* wizard, select the destination of the backup file from the **Backup to** drop-down list:
  - **My Computer:** Back up system files to a local directory on your computer. This is the default.
  - **TFTP:** Back up system files to a TFTP server.
    - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter **/archive/backups/Infoblox\_2009\_10\_20\_15\_30**. The directory path cannot contain spaces and backslash (\).
    - **IP Address of TFTP Server:** Enter the IP address of the TFTP server to which you want to back up the system files.
  - **FTP:** Back up system files to an FTP server.
    - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter **/archive/backups/Infoblox\_2009\_10\_20\_15\_30**. The directory path cannot contain spaces and backslash (\).
    - **IP Address of FTP Server:** The IP address of the FTP server.
    - **Username:** Enter the username of your FTP account.  
Note that if you have configured AD server for authentication, you must specify "domain name//username".
    - **Password:** Enter the password of your FTP account.
  - **SCP:** Back up system files to an SSH server that supports SCP.

- **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30`. The directory path cannot contain spaces and backslash (\).
- **IP Address of SCP Server:** The IP address of the SCP server.
- **Username:** Enter the user name of your SCP account.  
Note that if you have configured AD server for authentication, you must specify "domain name//username".
- **Use Keys:** If you select this checkbox, you can back up files to SCP without entering the password. The first time you select the checkbox, you need to enter the password. However, during subsequent times, the Infoblox server verifies whether Infoblox keys are available on the SCP server. If they are available, you can click the **Backup** button without entering the password. If Infoblox keys are not available on the SCP server, the following message is displayed:  

```
Infoblox SSH keys are not present on SCP server. Please upload the keys to the SCP server or download the keys and manually add it to the SCP server.
```
- **Password:** Enter the password of your SCP account.
- **Keys Type:** Select the SSH key type to be uploaded. At present, only ECDSA and RSA keys are supported. Click **Upload Keys** to upload the keys to the SCP server. If the keys are not available, click **Download Keys** to download the keys and manually add them to the SCP server.

Notes:

- If you are using Fedora, ECDSA keys are supported only on Fedora versions later than Fedora 12.
  - When you select **FTP** or **SCP**, ensure that you have a valid user name and password on the server prior to backing up the files. Also ensure that the target SSH server has the required permissions for an SCP backup. The permission must be 755 and the target server must have write permission to the directory to which you upload the backup file.
  - For an SCP backup, ensure that you are logged in as the user for whom the key was created. Also ensure that the `.ssh` directory on the server and the files it contains, have the correct permissions: 

```
chmod 600 ~/.ssh/authorized_keys && chmod 700 ~/.ssh/
```
  - If you promote a Grid Master or perform an HA failover, you must upload the SSH key once again for a successful SCP backup using keys.
3. If the Grid has a discovery member, Grid Manager displays the **NIOS data** and **Discovery data** checkboxes. You can select the **NIOS data** checkbox, to back up NIOS configuration data for the Grid and select the **Discovery data** checkbox, to back up discovery data for the Grid.  
If the Grid has a reporting member, Grid Manager displays the **Infoblox Reporting & Analytics App** checkbox. You can select the **Infoblox Reporting & Analytics App** checkbox, to back up Splunk application reporting data.
  4. Click **Backup**.

## Downloading Backup Files

You can save an existing backup file, or create and save a new one to your local management system, a TFTP server, an FTP server, or a SCP server.

To download an existing backup file:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Backup -> Manage Local Backup** from the Toolbar. Grid Manager displays the current backup files in the *Manage Local Backups* dialog box.
2. To download a backup file, select the checkbox of a backup file, and then click the Transfer icon. You cannot select multiple files for downloading.
3. Select one of the following from the **Backup to** drop-down list:
  - **My Computer:** Backup to a local directory on your computer. This is the default.
  - **TFTP:** Save the backup file to a TFTP server.
    - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter

```
/archive/backups/Infoblox_2009_10_20_15_30
```

- **IP Address of TFTP Server:** Enter the IP address of the TFTP server to which you want to save the backup file.
- **FTP:** Save the backup file to an FTP server.
  - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter

```
/archive/backups/Infoblox_2009_10_20_15_30.
```

- **IP Address of FTP Server:** The IP address of the FTP server.
- **Username:** Enter the username of your FTP server account.
- **Password:** Enter the password of your FTP server account.
- **SCP:** Save the backup file to an SSH server that supports SCP.
  - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30`.
  - **IPAddressofSCPServer:** The IP address of the SCP server.
  - **Username:** Enter the username of your SCP server account.
  - **Password:** Enter the password of your SCP server account.Note that when you select **FTP** or **SCP**, ensure that you have a valid username and password on the server prior to backing up the files.

4. Click **Transfer Copy**.

## Restoring Backup Files

You can restore a backup file of a NIOS configuration or a Discovery database to an appliance running the same NIOS version as that of the appliance from which the backup file originates. You can also restore a backup file from an appliance running a NIOS version to an appliance running a later NIOS version as long as the upgrade from the earlier NIOS version to the later version is supported. For example, you can restore a backup file from an appliance running NIOS 6.10.0 to an appliance running NIOS 7.3.200 because upgrading from NIOS 6.10.0 to 7.3.200 is supported. However, you cannot restore a backup file from an appliance running NIOS 6.9.0 to an appliance running NIOS 7.3.200 because upgrading from NIOS 6.9.0 to 7.3.200 is not supported.

You can restore an existing backup file on the appliance from which it originates, or restore a backup file from a different appliance (referred to as a forced restore). To download a backup file from a different appliance, see [Downloading Backup Files from a Different Appliance](#) below.

You must log in with a superuser account to back up and restore files. NIOS provides three ways to restore a backup file:

- From a local directory or the management system you use to operate the appliance
- From a TFTP server
- From a remote server using FTP. This option requires that you have a valid username and password on the FTP server prior to performing a backup or restore.



### Note

When you restore NIC interfaces to a VM, ensure that you provision appropriate NIC interfaces with the database content that must be restored to avoid any errors.

To restore a backup file to the same independent appliance or Grid Master:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Restore** from the Toolbar.
2. In the *Restore* dialog box, choose one of the following from the **Restore from** drop-down list:
  - **My Computer:** Restore a file from your local computer. This is the default.
    - **Filename:** Click **Select File** to navigate to the configuration file.
  - **TFTP:** Restore a file from a TFTP server.

- **Filename:** Enter the directory path and the file name you want to restore. For example, you can enter  

```

/archive/backups/Infoblox_2009_10_20_15_30

```
  - **IP Address of TFTP Server:** Enter the IP address of the TFTP server from which you restore the configuration file.
  - **FTP:** Restore a file from an FTP server.
    - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter  

```

/archive/backups/Infoblox_2009_10_20_15_30.

```
    - **IP Address of FTP Server:** The IP address of the FTP server.
    - **Username:** Enter the username of your FTP server account.
    - **Password:** Enter the password of your FTP server account.
    - **Grid Master (Local):** Restore from a local directory on the Grid Master. In the *Backup Set* table, select the file you want to restore.
3. To restore NIOS configuration data, select the **NIOS data** checkbox.
  4. To restore Discovery data, select the **Discovery data** checkbox. Discovery data should be restored to Consolidator appliances with the correct licensing.
  5. To download a backup file from one appliance to a different appliance, select **Force Restore from Different Grid** to enable the feature, and then select one of the following:
    - **Retain Current Grid Master IP Settings** (this is the default)
    - **Overwrite Grid Master IP Settings**
  6. Click **Restore**. In the *Confirm Restore* dialog box, click **Yes**.  
After restoring the file, the appliance restarts. The restore process overwrites all existing data. All pending scheduled tasks are not restored or reverted.
  7. Close your current browser window, wait a few minutes, and then reconnect to the NIOS appliance.



#### Note

You cannot restore a backup from a HA GM to a single node GM, if the HA GM has specific HA passive configuration.

## Downloading Backup Files from a Different Appliance

When you "force restore" a NIOS appliance, you download a backup file from one appliance to a different appliance. To restore a backup file to the same appliance or Grid Master, use the Restore function as described in Restoring Backup Files.

To download a backup file from one appliance to a different appliance:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Restore** from the Toolbar.
2. In the *Restore* wizard, do the following:
  - **Restore from:** Choose a source from which you restore the configuration file, as described in Restoring Backup Files.
3. Select **Force Restore from Different Grid** to enable the feature, and then select one of the following:
  - **Retain Current Grid Master IP Settings** (this is the default)
  - **Overwrite Grid Master IP Settings**
4. Click **Restore**. In the *Confirm Restore* dialog box, click **Yes**.  
After restoring the file, the appliance reboots. The restore process overwrites all existing data. All pending scheduled tasks are not restored or reverted.
5. Close your current browser window, wait a few minutes, and then reconnect to the NIOS appliance.



## Managing Upgrade Groups

To minimize the impact of Grid upgrades on your system operations, you can organize members into upgrade groups and schedule their software distributions. This is useful, for example, in a large Grid spanning multiple time zones where there are fluctuating network and downtime considerations at various locations. Note that you can also schedule their upgrades, depending on the existing releases and their upgrade paths. For information about the different upgrade methods, see [Lite Upgrades](#) and [Full Upgrades](#)

You can also import and export upgrade groups and their schedules in CSV format. For more information, see [Infoblox CSV Import Reference](#).

Infoblox provides two default upgrade groups:

- **Grid Master:** After you configure the Grid Master, it automatically becomes the only member of this group. You cannot modify or delete this group.
- **Reporting Member:** After you configure a reporting member in a Grid, it automatically becomes the only member of this group. This group will be upgraded automatically after the Grid Master and before other upgrade groups. You cannot modify, delete, or schedule this upgrade group. For information about reporting, see [Infoblox Reporting and Analytics](#).
- **Reporting Member (when you have configured reporting clustering Grid):** When you upgrade from NIOS 7.3.200 to a later release, the Grid will be upgraded to the single indexer mode. You can change the configuration after the upgrade. During an upgrade, the reporting members of the primary site are moved into separate groups to ensure that there are always some peers available to receive events sent by the forwarders. To achieve this, the replication factor must be equal to the search factor and the total number of reporting appliances. During an upgrade, data loss might happen if the peers that hold all searchable copies of a bucket are in the same group and they are all offline at the same time. To avoid this, ensure that you group the members in separate groups.
- **Default:** This is the default upgrade group to which the appliance automatically assigns Grid members. If you do not explicitly assign a member to an upgrade group, it remains in the Default group. You cannot delete or rename this group.  
Make sure that the default group has at least one member associated with it, otherwise the appliance displays that the upgrade process is still in progress even though it is complete. To avoid this, you can either use the `Infoblox > set grid_upgrade forced_end` command to stop the upgrade process or keep at least one member in the default group.

Grid Manager provides information about the upgrade group to which a member belongs. You can add or delete an upgrade group and monitor the software version that is currently running on the Grid and on individual member. You can do the following:

- Add an upgrade group, as described in [Adding Upgrade Groups](#) below.
- Modify an upgrade group, as described in [Modifying Upgrade Groups](#) below.
- View upgrade group information, as described in [Viewing Upgrade Groups](#) below.
- Delete an upgrade group, as described in [Deleting Upgrade Groups](#) below.

## Adding Upgrade Groups

When you create an upgrade group, you select the Grid members for that group, and specify whether the software distribution and upgrade occur on all group members at the same time, or successively in the order they are listed in the group members list. A Grid member can belong to only one upgrade group.





## Note

The appliance displays a warning message when you create an upgrade group that includes the two peers of a DHCP failover association. Infoblox recommends that you assign DHCP failover peers to separate upgrade groups to minimize the risk of a loss in DHCP services. For example, if DHCP failover peers are in the same upgrade group and its members upgrade simultaneously, the upgrade causes a loss in DHCP services.

Note the following recommendations when you create an upgrade group:

- Put the following members in the first upgrade group after the Grid Master upgrade: all Grid Master candidates, DNS primaries, and the DHCP logging member.
- To minimize the risk of a loss in DNS services, put the name servers for a zone in different upgrade groups, and assign the primary and secondary servers to separate upgrade groups.

To add an upgrade group:

1. From the **Grid** tab, select the **Upgrade** tab.
2. Click **Toggle Group List View** to display the list of upgrade groups, and then click the Add icon.
3. In the *Add Upgrade Group* wizard, complete the following:
  - **Name:** Enter a name for the upgrade group. The name can contain alphanumeric characters, spaces, underscores, hyphens, and dashes.
  - **Distribute to Members:** Select one of the following to specify how the Grid Master distributes software to the members in the group.
    - **Simultaneously:** Select this to distribute software upgrade files to all group members at the same time.
    - **Sequentially:** Select this to distribute software upgrade files to group members in the order they are listed in the group members list.
  - **Upgrade Members:** Select one of the following to specify how the group members upgrade to the new software version.
    - **Simultaneously:** Select this to upgrade all group members at the same time.
    - **Sequentially:** Select this to upgrade group members in the order they are listed in the group members list.
  - **Comment:** Enter useful information about the upgrade group, such as the location of the group.
4. Click **Next** to select members for the group. Complete the following:
  - Click the Add icon. Grid Manager adds a row to the Member Assignment table.
  - Click **Select**. In the *Member Selector* dialog box, select the members you want to add to the group, and then click the Select icon. Use Shift+click and Ctrl+click to select multiple members. Note that if you choose to distribute and upgrade members sequentially, the distribution and upgrade occur in the order the members are listed. You can reorder the list by dragging a member to a desired location or by selecting a member and using the up and down arrows next to the checkbox to place the member at a desired location. You can also delete a member from the list.

After you add a member, the appliance adds it to the group members list. The first Grid member in the list determines the time zone of the group when you schedule the distribution and upgrade. Therefore, Grid Manager displays the time zone of the first Grid member in the list. (For information about setting time zones, see [Managing Time Settings](#).)
5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Modifying Upgrade Groups

You can modify an existing upgrade group to change the group name or how the distribution and upgrade are performed. You can also add and delete members.

To modify an upgrade group:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Toggle Group List View**.
2. Select an *upgrade\_group* checkbox, and then click the Edit icon in the row. You can also click the Edit icon directly without selecting the checkbox.
3. The *Upgrade Group* editor provides the following tabs from which you can modify data:

- **General:** Modify the fields as described in Adding Upgrade Groups.
  - **Member Assignment:** Add or delete members as described in Adding Upgrade Groups.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Viewing Upgrade Groups

In the **Upgrade** tab, Grid Manager lists the Grid Master group, the Default group, and other upgrade groups you have configured. You cannot modify or delete the Grid Master group. You can modify the Default group, but you cannot delete it. To view the members in a specific upgrade group, click the arrow next to the group name to expand the group. All groups are collapsed by default.

Before a distribution or upgrade starts, you can move members from one group to another, reorder the members, or remove a member from an upgrade group. The member you remove automatically joins the Default group. (For information, see [Managing Distributions](#), in *Upgrading NIOS Software*.) You cannot add, delete, or reorder members in an upgrade group while a distribution or upgrade is in progress. You can skip a member in an upgrade group from a distribution only before the distribution starts, or after you pause it. For information, see [Pausing and Resuming Distributions](#) in *Upgrading NIOS Software*.

To view the upgrade groups in a Grid:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Toggle Group List View**.  
Grid Manager displays the Grid Master at the top of the list. All other upgrade groups are listed alphabetically after the Grid Master. You can click the arrow next to a group to view members in the group.
2. Grid Manager displays the following:
  - **Group:** The name of an upgrade group to which the member belongs.
  - **Member:** The name of the member.
  - **Status:** Displays the overall status of an upgrade group at the group level and individual status for each member when you expand the upgrade group. At the group level, this displays the most severe status among the members. For example, when there are three out of five members are offline, the overall status shows **3 of 5 members** in red, which means offline.
  - **IP Address:** The IP address of the member.
  - **Running Version:** The NIOS software version that is currently running on the member.
  - **Distribution Status:** The distribution status of the group.
  - **Timestamp:** The date, time, and time zone when a distribution or upgrade is complete.

You can hide some of the default columns, but you cannot sort the information in this table. You can use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches. You can also create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#), in *Finding and Restoring Data*.

## Deleting Upgrade Groups

When you delete an upgrade group, members in the upgrade group that you want to delete will be moved to the Default group. Grid Manager displays a warning before deleting an upgrade group.

To delete an upgrade group:

1. From the **Grid** tab, select the **Upgrade** tab, and then click **Toggle Group List View**.
2. Select an *upgrade-group* checkbox, and then click the Delete icon.
3. In the *Delete Confirmation* dialog box, click **Yes**.

## Full Upgrades

A full upgrade occurs when there are database schema changes between the existing and upgrade releases. In general, when you upgrade to a major release, you are performing a full upgrade. Depending on the upgrade and revert paths that your existing release supports, you may or may not be able to schedule a full upgrade. A full upgrade that cannot be scheduled does not allow for data replication between the Grid Master and members. For information about supported upgrade and revert paths, refer to the latest release notes on the Infoblox Support site.

Depending on the upgrade paths your current release supports, when you schedule a full upgrade, the Grid Master

immediately replicates certain core network service tasks to Grid members while putting other tasks in queue until the members have been upgraded. For information about which data and tasks the Grid Master replicates to members immediately, see [Guidelines for Scheduling Full Upgrades](#) . For information about how to schedule an upgrade, see Scheduling Upgrades in [Upgrading NIOS Software](#).

# Using the Grid Manager Interface

This section explains the different components of Grid Manager and details how to use each of them. It covers the following topics:

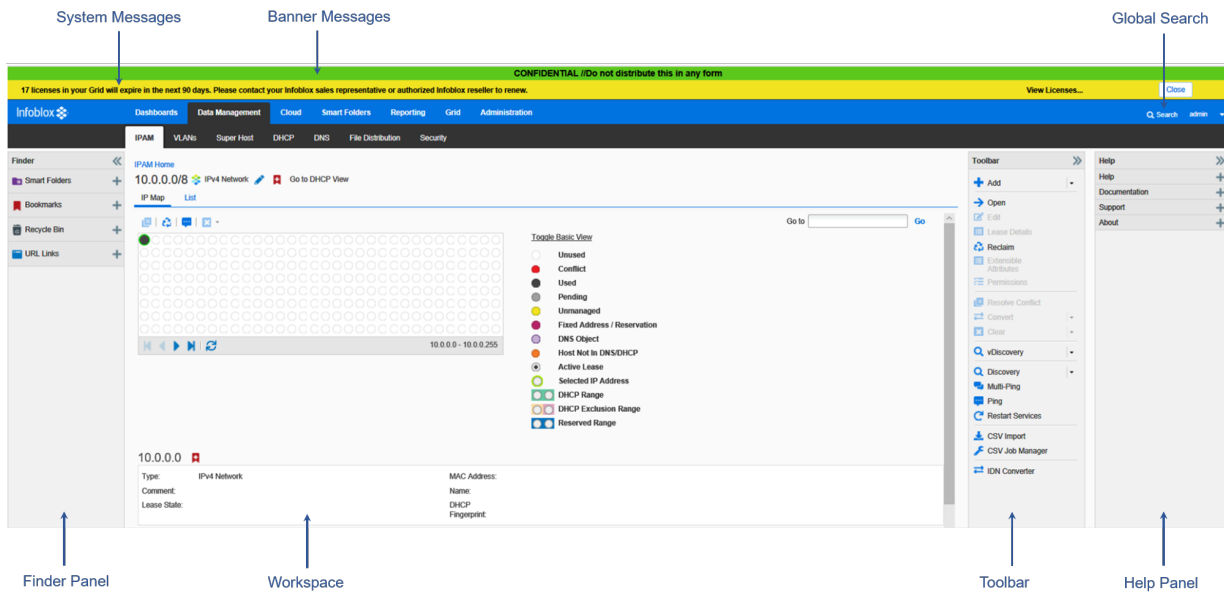
- [About the Grid Manager Interface](#)
- [Scheduling Tasks](#)
- [Scheduling New IPAM/DHCP Objects and Associated Port Configurations](#)
- [Performing Basic Operations Using Grid Manager](#)
- [Dashboards](#)
- [Smart Folders](#)
- [Importing and Exporting Data using CSV Import](#)
- [SSL and TLS Protocols](#)
- [Managing Certificates](#)
- [Configuring Approval Workflows](#)

## About the Grid Manager Interface

Grid Manager provides an easy-to-use interface that simplifies core network services management. Its navigational tools enable you to quickly move through the application and retrieve the information you need. You can customize different elements in your workspace, and hide and display panels as you need them. It also provides different types of Help, so you can immediately access the information you need to complete your tasks.

The following Grid Manager Interface figure illustrates the typical layout of Grid Manager. It identifies common elements of the interface and features that you can use:

*Grid Manager Interface*



## System Messages

Grid Manager displays system messages at the top of the screen. In wizards and editors, it displays messages at the top as well.



### Note

Some configuration changes require a service restart. Grid Manager displays a message whenever you make such a change. Click the **Restart** icon that appears in the message to restart services.

## Security and Informational Banners

Grid Manager displays banner messages on the header and footer of the screen. Only superusers can publish the informational and security banner. There are two types of banners:

- Security Banner - Security banner indicates the security level of the Infoblox Grid. There are five security levels to choose from the Security list box. The available security levels are Top Secret, Secret, Confidential, Restricted, and Unclassified.
- Informational Banner - You can use the informational banner for multiple uses, such as to indicate whether the Infoblox Grid is in production or a lab system. You can also publish messages of the day.

For more information about configuring security level banner and configuring informational level banner, see [Managing a Grid](#).

## Breadcrumbs Navigation

Breadcrumbs navigation displays your path to the current page. It helps you keep track of your location in Grid Manager. You can click any of the links to get back to a previous page.

## Global Search

Use Global Search to find data. Grid Manager searches the entire NIOS database for data that matches the criteria you specify. For additional information on Global Search, see [Finding and Restoring Data](#).

## Finder Panel

The *Finder* panel appears on all pages in Grid Manager. It provides the following tools:

- Smart Folders: Use smart folders to organize your data according to criteria that you specify.
- Bookmarks: Stores data that you have marked for easy retrieval.
- Recycle Bin: Stores deleted objects that you can either restore or permanently remove.
- URL Links: You can add, modify, and delete third party URL links of frequently used portals and destination pages.

You can resize, collapse, and expand the Finder panel.

## Toolbar Panel

The vertical Toolbar panel provides easy access to commands. The Toolbar is available in all pages, except the Dashboard. Its content changes depending on the type of data displayed in the work area. You can resize, collapse, and expand the *Toolbar* panel.

## Help Panel

The *Help* panel provides the following types of Help:

- **Help:** Expand this section to view information about the window currently displayed.
- **Documentation:** Expand this section to download the latest versions of the Infoblox Administrator Guide and Infoblox API Documentation.
- **Support:** Expand this section to view links to the Infoblox web site and Technical Support site.
- **About:** Expand this section to view information about the NIOS software version. You can resize, collapse, and expand the Help panel. In addition, each dialog box also provides a Help panel that contains information specific to the dialog box. You can expand and collapse the Help panel in dialog boxes as well.

## Wizards and Editors

Grid Manager provides a wizard for every object that you can create. You use wizards to enter basic information required to create an object. If you want to configure additional parameters, you can then save the object and edit it.

Note that all required fields are denoted by asterisks.

Your connection to Grid Manager may time out if a save operation takes longer than 120 seconds to complete. This can occur when multiple, complex operations are initiated by several users. It does not result in any data loss.

## Tooltips

Tooltips display the function of each button. Hover your mouse over a button or icon to display its label.

## Customizing Tables

Grid Manager uses dynamic tables to display information. You can customize tables by resizing columns, sorting the data, and selecting certain columns for display. Your settings remain active until you log out.

To resize columns in a table:

1. In the table, place your pointer on the right border of the header of the column you want to resize.
2. Drag the border to the desired width.

To sort the data displayed in a table, click the header title. You can click the header title again to reverse the sort order. Alternatively, you can do the following:

1. In the table, mouse over to a header title and click the down arrow key.
2. Select **Sort Ascending** or **Sort Descending**. To edit columns:
3. In the table, mouse over to a header title and click the down arrow key.
4. Select **Columns > Edit Columns**.
5. Do the following:
  - **Width:** Specify the width of the column in pixels. The minimum is five and the maximum is 999.
  - **Sorted:** Indicates whether the data in the column can be sorted
  - **Visible:** Click the checkboxes of the columns you want to display, and clear the checkboxes of those you want to hide.
6. Do one of the following:
  - Click **Apply** to apply your settings to the column.
  - Click **Cancel** to close the editor without saving your settings.
  - Click **Reset** to reset the settings to the default.

Grid Manager displays the selected column in the table.

To reorder columns in a table, drag and drop the columns to the desired positions.

## Selecting Objects in Tables

In a table, Grid Manager displays data on multiple pages when the number of items to be displayed exceeds the maximum number of items that can be displayed on one page. Use the navigational buttons at the bottom of the table to page through the display.

You can select multiple rows in a table. For example, in a Windows browser, you can do the following to select multiple rows:

- Use SHIFT+click to select multiple contiguous rows.
- Use CTRL+click to select multiple non-contiguous rows.
- Click the checkbox in the table header to select all rows on a page, as shown in the below figure Select all in a Table.

When you click the select all checkbox in a table that contains multiple pages, only the rows on the current page are selected. Grid Manager displays a message that indicates the total number of selected rows on the page. You can click **Select all objects in the dataset** to select all rows in the entire table. When you select all rows in the table, Grid Manager displays a message to indicate that. You can then click **Clear Selection** to deselect the rows.

After you select all rows on a page, you can deselect a specific row by clearing the checkbox of the row. You can also click a row (not the checkbox) in the table to select the item and deselect the others.

In a table, when you select all the objects for deletion, the objects that are not deleted from the database remain in the table after the operation is completed.

### Select All in a Table

Annotations in the screenshot:

- Click this link to select all rows on all pages (points to "Select all objects in this dataset")
- Click this check box to select all rows on this page only (points to the "Select all" checkbox in the table header)
- Use these navigational buttons to page through the display (points to the table's navigation controls)

Network	Comment	IPAM Utilization	Discovery Engine	Discovered ...	Discovered VLA ...	Assigned VLAN ID	Assigned VLAN ...	VRF Name
10.150.0.0/16	Park Place	75.0%	None					
10.151.0.0/16	Krakow, Poland	75.0%	None					
10.152.0.0/16	Sofia, Bulgaria	75.0%	None					
10.153.0.0/16	Luxembourg	75.0%	None					
10.199.0.0/16	Sydney	75.0%	None					
10.200.0.0/16	Singapore	74.0%	None					
10.253.0.0/16	Bulgaria	55.0%	None					
10.254.0.0/16	Office in a Box	1.0%	None					
23.21.100.0/24	Internet Facing	0.3%	None					
23.21.200.0/24	Internet Facing	0.3%	None					
23.21.206.0/24	Internet Facing	0.3%	None					

## Modifying Data in Tables

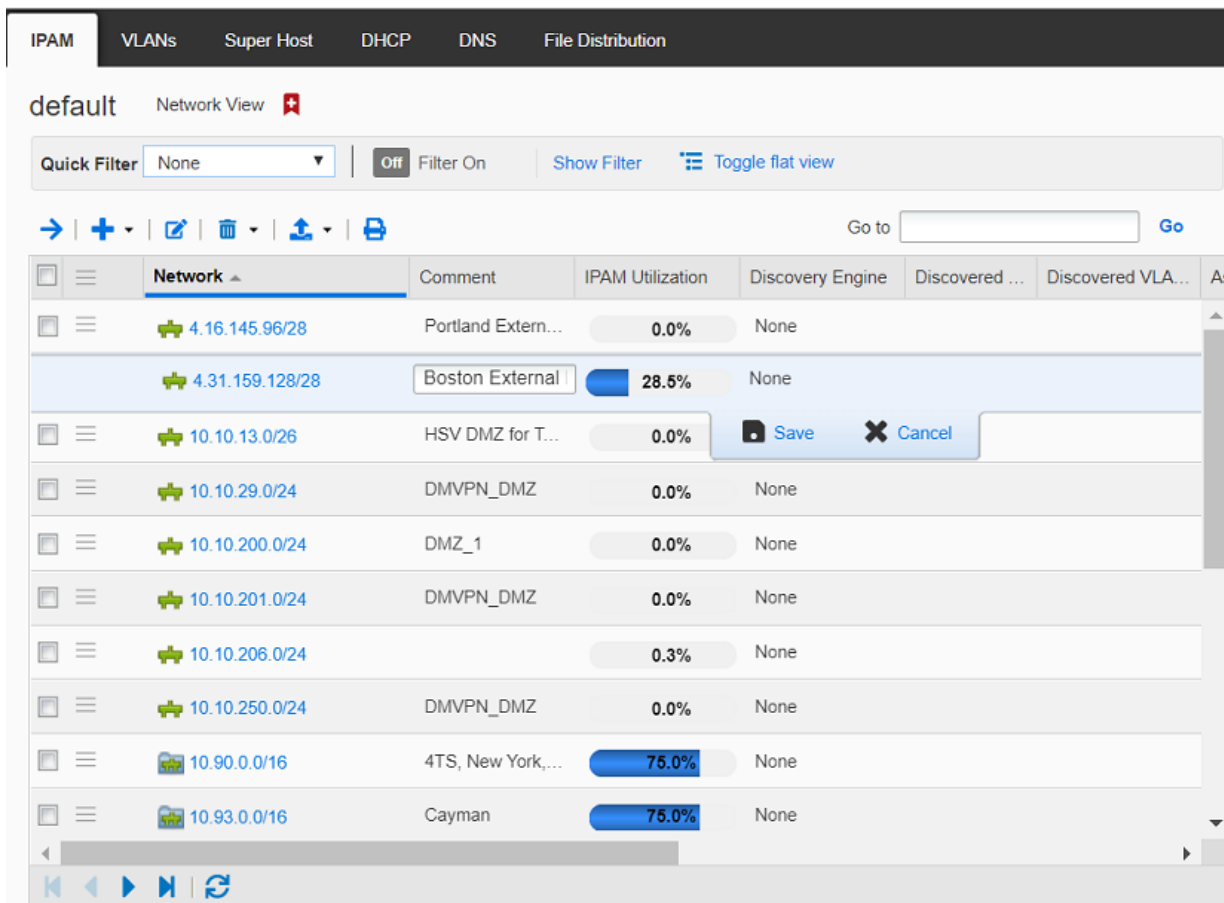
Infoblox provides inline editing for certain fields in some tables. You can use this feature to modify data directly in a table instead of going through an editor.

To update information in a table, you must have read/write permission to the data. When you enter or select a new value, the appliance validates the data format before saving the updated data.

To modify data in a table:

1. From any panel that supports inline editing, double click the row of data that you want to modify. The appliance displays the inline editing editor in the selected row, as shown in the Inline Editing figure.
2. Depending on the data type, enter the new data in the field or select an item from the drop-down list. Note that some fields are read-only.
3. Click **Save** to save the changes, or click **Cancel** to discard them.

### Inline Editing



## Printing from Grid Manager

In Grid Manager, you can print information from panels and pages that support the Print function. Grid Manager prints data one page at a time. The amount of data that is displayed in a specific panel depends on the table size configuration that you set in your user profile. For more information about specifying the table size, see [Setting Login Options](#). To print:

1. Click the **Print** icon. You must allow pop-up windows in your browser for printing. Grid Manager displays a separate browser window.
2. Click **Print**.  
Grid Manager displays the *Print* dialog box.
3. Configure printer settings and parameters.
4. Depending on your browser, click **OK** or **Print**.

## Scheduling Tasks

When you perform a task, such as adding a DNS zone or modifying a DHCP range, you can execute it immediately or schedule it for a future date and time, depending on your permissions. The scheduling feature is useful when you want to add, modify, or delete a record, or schedule a network discovery at a desired date and time. Using this feature, you can streamline your day-to-day operations. For example, you can schedule the deletion of records that you use for testing when the test time is up. You can also reassign an IP address to a fixed address when the location of the server to which the fixed address is assigned changes from one network to another.



Certain tasks, scheduled or not, may be subject to approvals if approval workflows are defined for specific admin groups. For information about how to define submitters and approvers for an approval workflow, see [Configuring Approval Workflows](#).

Note that not all tasks can be scheduled or routed for approval. For a list of supported objects, see [Supported Objects for Scheduled and Approval Tasks](#) below.

When you schedule a task or submit it for approval, consider the following:

- The appliance cannot execute a scheduled or approval task that is associated with an extensible attribute, if you delete the extensible attribute after you have scheduled the task or submitted it for approval. For information about extensible attributes, see [About Extensible Attributes](#).
- The appliance cannot execute, reschedule, or delete a task that is associated with a child object (such as a DHCP range) if you delete the parent object (such as a network) after you have scheduled the task or submitted it for approval.
- There are certain guidelines about scheduled and approval tasks when you upgrade the software, back up the database, and restore data. For information, see [Guidelines for Upgrading, Backing Up, and Restoring Data](#).

You can schedule the addition, modification, and deletion of certain objects. For a list of supported objects, see [Supported Objects for Scheduled and Approval Tasks](#) below.

Depending on your permissions and the admin group to which you belong, your scheduled tasks may be subject to approvals by other admins in your organization. You may or may not receive email notifications about the status of your scheduled tasks depending on the configuration of the approval workflows. Approvers can reschedule your tasks after they have approved the tasks, if they have scheduling permissions. When you schedule and submit a task, you may need to enter a ticket number associated with the task or a comment about the task. For more information about approval workflows, see [Configuring Approval Workflows](#).

Only superusers can view, reschedule, and delete all scheduled tasks. Limited-access admins can reschedule and delete only their scheduled tasks. If your scheduled tasks require approvals, the approvers who have scheduling permissions may reschedule your tasks to a different date and time after they have approved the tasks. Depending on your admin permissions, there are certain scheduled and approval tasks that you may or may not be able to perform. For more information about supported tasks for different admin groups, see [Configuring Approval Workflows](#).

The appliance sends email notifications to local admins, except for those who do not have email addresses, when email notification is enabled for the admins and any of the following happens:

- A superuser schedules a task, and another superuser reschedules or deletes the task.
- A limited-access admin schedules a task, and a superuser reschedules or deletes the task.
- A superuser or a limited-access admin schedules a task, and the task fails.
- An admin is configured to receive notifications based on the configuration of an approval workflow. For information about approval workflows, see [Configuring Approval Workflows](#).

Superusers can also grant scheduling permissions to other admin groups. When the scheduling permission is added or inherited from an admin role, limited-access admin groups can schedule tasks. For more information, see [Administrative Permissions for Discovery](#).

## Supported Objects for Scheduled and Approval Tasks

- DNS zones (authoritative, forward, stub, and delegated)
- DNS views
- DNS resource records (except SOA records)
- Import resource records to DNS zones
- Lock and unlock DNS zones
- Hosts
- Bulk hosts
- Roaming hosts
- Shared records
- Shared record groups
- IPv4 and IPv6 networks
- IPv4 and IPv6 network containers
- IPv4 and IPv6 shared networks

- IPv4 and IPv6 DHCP ranges
- IPv4 and IPv6 reserved ranges
- IPv4 and IPv6 fixed addresses
- IPv4 reservations
- DHCP fingerprints
- DHCP filters (MAC, option, NAC, relay agent, and DHCP Fingerprint)

#### Note

Only IPv4 MAC filters support approval workflows.

- IPv4 MAC address filter items
- Conversion of IPv4 and IPv6 static and dynamic leases
- Microsoft objects that are supported by NIOS
- Load balancer related objects
- DNS64 Synthesis Groups
- All IPAM tasks except CSV imports
- Response Policy Zones
- Response Policy records
- VLANs

You can also schedule the following operations or create approval workflows for them:

- Network Discovery
- Device Discovery
- VM Discovery
- Port Control provisioning tasks for setting Admin Status, VLAN assignments, and a Description;
- Defining infrastructure device port reservations for the following IPAM/DHCP objects:
  - IPv4 Reservations
  - IPv4/IPv6 Fixed Addresses
  - Hosts
  - Grid Members
  - IPv4 and IPv6 Networks
- Service restarts (for scheduled tasks only)

#### Note

Service restarts are not subject to approvals.

## About Long Running Tasks

A long running task is a task that requires more than 30 seconds to complete and involves a large amount of data. When Grid Manager performs a long running task, it displays the *Long Running Task* dialog box that indicates whether you can run the task in the background. You can navigate to another tab or perform other functions only if the task can be run in the background. For information, see [Running Tasks in the Background](#) below.

Grid Manager disconnects if a task takes more than five hours to perform. Though you can log back in to Grid Manager while the appliance continues to perform the task, Grid Manager does not display the progress of the task.



#### Note

You cannot stop a long running task once you start it.

The appliance supports the following long running tasks:

- Restoring the database
- Backing up the database
- Backing up licenses

- Signing DNS zones
- Unsigning DNS zones
- Exporting DS records and trust anchors
- Deleting all objects in a table or dataset
- Modifying multiple extensible attributes
- Viewing DNS and DHCP configuration properties
- Migrating bloxTools data
- IPAM tasks on the Tasks Dashboard
- Downloading the following:
  - Audit logs
  - Syslog files
  - Support bundles
  - SNMP MIB files
  - NTP keys
  - HTTPS certificates
  - Traffic capture

## Running Tasks in the Background

Grid Manager allows certain long running tasks to run in the background. You can navigate to other tabs and perform other functions when Grid Manager performs tasks in the background. However, when you make changes to objects that are currently affected by a long running background task, Grid Manger does not save the changes until after the long running task is completed. Grid Manager can perform up to 10 background tasks at a time.

You can run the following tasks in the background:

- Signing DNS zones
- Unsigning DNS zones
- Modifying multiple extensible attributes
- Deleting all objects in a table or dataset
- Migrating bloxTools data

To run a task in the background:

1. Perform the task following the instructions described in this guide.
2. In the *Long Running Task* dialog box, click **Run in Background**.

You can view the progress of the task by clicking the progress bar at the top of the interface.

## Monitoring Long Running Tasks

When you have one or more tasks running in the background, Grid Manager displays a progress bar next to the Global Search icon at the top of the interface. You can click the progress bar to view detailed information about the tasks in the *Background / Long Running Task* viewer. In this viewer, Grid Manager displays a progress bar for each task that is currently running in the background. When all background tasks are completed, the progress bar at the top of the interface disappears. Grid Manager displays a message at the top of the interface when the task is completed successfully or if the task fails.

For other tasks that you cannot run in the background, the *Long Running Task* dialog box remains open until the task is completed. You cannot navigate to other tabs or perform other functions when the long running task is in progress. Grid Manager closes the dialog box when the task is completed. It also displays a message at the top of the interface when the task is completed successfully or if the task fails.

## Viewing Tasks

The appliance displays scheduled tasks and approval tasks in the **Task Manager** tab of Grid Manager. Scheduled tasks are those with scheduled time listed and approval tasks contain approval status. A task can also be scheduled and queued for approval at the same time. By default, all completed and rejected tasks are displayed in **Task Manager** for up to 14 days before they are removed from the list. You can configure how long the completed and rejected tasks are displayed in **Task Manager** using the CLI command `set delete_tasks_interval`. For more information about the CLI

command, see *Infoblox CLI Guide*.

The appliance logs all tasks in the audit log and associates each with a task ID. By default, Grid Manager sorts tasks by Task ID in **Task Manager**. You can view tasks that you are allowed to see based on your permissions. For information about admin permissions, see [About Administrative Permissions](#).

To view tasks:

1. From the **Administration** tab, select the **Workflow** tab -> **Task Manager** tab.
2. Grid Manager displays the following information for each task:

You can do the following in the **Task Manager** tab:

Field Name	Description
<b>Action</b>	The operation the appliance performs in this task. The can be: <b>Add, Modify, Delete, Network Discovery, Lock/Unlock Zone, or Restart Services</b> .
<b>Affected Object</b>	The name or value of the object that is associated with the task. For example, if the task involves an A record, this field displays the domain name of the record. If it is a fixed address, it displays the IP address of the fixed address.
<b>Approval Status</b>	The current approval status. Possible values are <b>Approved, Not Applicable, Pending, and Rejected</b> .
<b>Approver</b>	The username of the admin who has approved this task.
<b>Approver Comment</b>	Comments entered by the approver.
<b>Associated Task</b>	(hidden by default) Applies to Port Configuration tasks. If a port configuration task is dependent on an object change task (such as a new Fixed Address or an edit to an existing object), this could will show the Task ID value for the associated object change task.
<b>Executed Time</b>	The date, time, and time zone the task was executed.
<b>Execution Status</b>	The execution status of the task. Possible values are Completed, Failed, Pending, and Executing.
<b>Object Type</b>	The object type. For example, the appliance can display A Record or Fixed Address.
<b>Scheduled Time</b>	The date, time, and time zone when the appliance executes the task.
<b>Submitted Time</b>	The date, time, and time zone when the task was submitted.
<b>Submitter</b>	The username of the admin who scheduled or submitted the task.
<b>Submitter Comment</b>	Comments entered by the submitter.
<b>Task ID</b>	The ID associated with the task. The appliance assigns an ID to a task in chronological order. By default, the appliance sorts tasks by Task ID.

Field Name	Description
<b>Task Details</b>	Detailed information about the task. This message also appears in the audit log.
<b>Ticket Number</b>	For an approval workflow, this number may be entered by the submitter to associate the task with a help desk ticket number or a reference number.
<b>Type</b>	Indicates key information about certain types of executing/executed jobs. The Type column lists values for Port Control and for Object Change tasks undertaken by Grid Manager or submitted by Grid Manager for approval by the administrator.

- Sort the tasks in ascending or descending order by column, except for **TaskDetails**.
- Use filters and the search function to look for specific values.



### Note

You cannot use the search function to search for approval or execution status. Use filters to search for these values.

- Create a quick filter to save frequently used filter criteria. Grid Manager provides the following default quick filters that you can select from the Quick Filter drop-down list: **PendingApprovals**, **RejectedTasks**, and **ScheduledTasks**. For more information about using quick filters, see [Finding and Restoring Data](#).
- Export and print the information in the table.
- Control the display of information in the panel by toggling between a single-line view and a multi-line view.
- Reschedule a task, cancel a scheduled task, or execute a task immediately.
- For approvers, select a task and click the Approve icon to approve the task, or click the Reject icon to disapprove the task. You can also reschedule the task while approving it.



### Note

If you have multiple pages of tasks in **Task Manager**, you can select multiple tasks on the current page for approval or disapproval. If you click the **Select all objects in this dataset** link to select all the tasks in the dataset, the Approve and Reject icons are disabled and you cannot approve or reject any task.

## Using the Task Manager Action Menu



The Task Manager page provides an Action icon  column with a series of menu options for features related to grid Manager tasks to manage task execution, scheduling and approval. Menu choices change based upon the context and the current state of tasks in the table; features available in the Action menu include the following:

Figure 1.6 Task Manager Action menu



Features in Action Menu	Description
( <i>Applies only with Network Insight</i> ) <b>Approve</b> and <b>Reject</b>	Enables admins to approve or reject a pending job; rejecting a job immediately cancels it.
<b>Associated Task (applies only with port configuration tasks)</b>	Choosing this option opens the object change task, if any, for the currently selected port configuration task.
<b>ExecutionLog</b>	Opens a completed task's execution log window. the Execution Log lists the complete communications sequence sent to a device to perform a port control task.
<b>ExecuteNow</b>	Force a selected pending task to execute immediately.
<b>Re-execute</b>	Allows you to re-run the selected task. Combined with the Execution Log, this process can aid in troubleshooting a failed port control task.
<b>Reschedule</b>	Opens the Reschedule window for the selected task. To immediately execute this task, click <b>Now</b> . Or, in the <i>Reschedule</i> panel, click <b>Later</b> , and then specify a date, time, and time zone. You can reschedule the task if you have the applicable permissions. Click <b>Save</b> to commit the changes.
<b>Delete</b>	Deletes the pending task.
<b>View</b>	Opens the Task Viewer to the currently selected task. For related information about using the Task Viewer to View Job Logs and Approve Jobs, see <a href="#">NetMRI Task Pack</a> .

## Scheduling New IPAM/DHCP Objects and Associated Port Configurations

Specific IPAM/DHCP object types support scheduling of Port Configuration tasks as part of their object definition:

- New IPv4 and IPv6 Networks;
- IPv4 Reservations;
- Fixed IPv4 Addresses;
- Fixed IPv6 Addresses;
- Hosts;
- Infoblox Grid Members (including HA Pairs; discussed in the sections beginning in [Adding Grid Members](#)).

For all of these object types, you can click a **Schedule for Later** button at any step in the Wizard to commit the currently defined settings for the new object to a new Grid Manager task. For example, consider creating a new IPv4 network. After defining the IP address for the new network, you simply wish to create it and not to define any further configurations. Click **Schedule for Later** to skip the intervening Wizard steps and display the final Scheduling page of the Wizard, an example of which appears in the following Create Network Schedule for Later page figure.

*Create Network Schedule for Later page*

1. To create the new object immediately, select **Now** and click **Save & Close**.
2. You can choose to have Grid Manager create the object at a later time. To do so, select **Later**. Choose a **Selected time** by entering or selecting a **Start Date** (click the calendar icon to choose a calendar date) and a **Start Time** (click the clock icon to choose a specific time in fifteen minute increments), and choose a **Time Zone**.

When you step through the entire Wizard process (without clicking **Schedule for Later**) and if you also define port configuration settings, the object creation wizards provide a final Scheduling page with separate scheduling definitions, for the object and for the object's port configuration, as shown in the following figure.

*Scheduling Create Object and Port Configuration*

The final step for creating the Network, Fixed Address, Host, or IP reservation is to define when the task executes, including associated port control definitions. The port configuration is performed in a separate task, defined on the same wizard page. The port configuration task can be done at the same time that Grid Manager provisions the object, or may be scheduled for a later time.

1. To create the new object immediately, select **Now**.

By selecting Now, no task is created by Grid Manager and it simply creates the object. The completed object creation appears in the audit log. For more information, see [Viewing Tasks](#). Grid Manager creates a task for object creation only when you use **Schedule for Later**. Also, all port configuration and network provisioning instances create a new task under the Task Manager.

2. You can instruct Grid Manager to create the network at a later time. To do so, select **Later**. Choose a **Selected time** by entering or selecting a **Start Date** (click the calendar icon to choose a calendar date) and a **Start Time** (click the clock icon to choose a specific time in fifteen minute increments), and choose a **Time Zone**.

3. Port configuration and network provisioning tasks can be synchronized to take place at the same time as the creation of the new object under IPAM/DHCP; if so, keep the **At same time as above** option. Scheduling Recursive Deletions

- a. You can also schedule the task at a different time. To do so, select **Later** (under **Port Configuration**). Choose a **Selected time** by entering or selecting a **Start Date** (click the calendar icon to choose a calendar date) and a **Start Time** (click the clock icon to choose a specific time in fifteen minute increments), and choose a **Time Zone**.

The Port Configuration provision cannot take place at a schedule time or date before the associated object creation.

Object creation must successfully complete before its associated port configuration task begins.

4. Click **Save & Close** to complete the configuration.





## Note

Port configuration tasks (or any operation that queries for or changes device configurations through Grid Manager, including Discovery) are subject to a feature called Blackout Periods, which are defined elsewhere in Grid Manager. Blackouts are a scheduled feature that instructs Grid Manager not to perform Discovery operations and Port Control/provisioning operations on the managed network at specified times, days and dates. See the section [Defining Blackout Periods](#) for details.

## Scheduling Creation of new IPv4/IPv6 Networks and Associated Device Provisioning

Defining new IP networks in Grid Manager supports scheduled Network Provisioning tasks on discovered and managed infrastructure devices as part of their object definition:

- Provisioning IPv4 networks
- Provisioning IPv6 networks

These networks are managed on the IPAM and DHCP pages (unless the network is excluded from DHCP), with the difference that Network Insight provisions the network directly on the specified router or switch-router. You may define new networks under IPAM and under DHCP as one scheduled task, and may provision those networks on devices that are discovered and managed under Network Insight, as another scheduled task.

Both tasks may be separately scheduled. They can take place immediately: the network is created first, then the network is provisioned on the device, along with the necessary device configuration, which Grid Manager also handles using the required CLI credentials; or either or both tasks may be scheduled for a later time.

The scheduled IPv4 or IPv6 network must be created under IPAM or DHCP in Grid Manager before scheduled device configuration or provisioning of networks on those devices can take place.

When you define a new network under IPAM or DHCP, you can click **Schedule for Later** in any Wizard page to skip further configuration in the Wizard and commit the network settings to Grid Manager. If you do not want to provision the network, clicking **Schedule for Later** will display a shorter scheduling page:




*Create Network Schedule page (After clicking Schedule for Later)*

Add IPv4 Network Wizard > Step 8 of 8

**Create IPv4 Network**

Now  
 Later

**Selected time:**

**Start Date**      2020-04-23   
**Start Time**      04:06:28 PM   
**Time Zone**      (UTC + 5:30) Bombay, India 

**Your time:**

2020-04-23  
 04:07:22 PM  
 (UTC + 5:30) Bombay, Calcutta,  
 Madras, New Delhi

In cases of this type, you schedule or execute only a single task: creating the new network under DHCP/IPAM. No network provisioning task takes place.

When you provision the network without clicking **Schedule for Later**, the wizard provides a final Scheduling page with an expanded set of two task schedules as shown in the below figure:

## Scheduling Add Network and Network Provisioning Tasks

Add IPv4 Network Wizard > Step 9 of 9

Now  
 Later

Selected time:

Start Date: 2020-03-26

Start Time: 10:09:55 AM

Time Zone: (UTC + 5:30) Bombay, C

Your time:

2020-03-26  
10:24:43 AM  
(UTC + 5:30) Bombay, Calcutta, Madras, New Delhi

Network Provisioning

Location: Eng. Lab. HQ

At same time as above  
 Later

Selected time:

Start Date: 2020-03-26

Start Time: 10:09:55 AM

Time Zone: (UTC + 5:30) Bombay, C

Cancel Previous Next Save & Close

1. To immediately create the new network, you can select **Now**.
2. You can choose to have Grid Manager create the network at a later time. To do so, select **Later**. Choose a **Selected time** by entering or selecting a **Start Date** (click the calendar icon to choose a calendar date) and a **Start Time** (click the clock icon to choose a specific time in fifteen minute increments), and choose a **Time Zone**.
3. The Network Provisioning task can be synchronized to take place at the same time as the creation of the new network under IPAM/DHCP; if so, keep the **At same time as above** option.
4. You can also provision the network onto the device at a differently scheduled time. To do so, select **Later** (under **Network Provisioning**). Choose a **Selected time** by entering or selecting a **Start Date** (click the calendar icon to choose a calendar date) and a **Start Time** (click the clock icon to choose a specific time in fifteen minute increments), and choose a **Time Zone**.



### Note

Network provisioning tasks (or any operation that queries for or changes device configurations through Grid Manager, including Discovery) are subject to a feature called Blackout periods, which are defined elsewhere in Grid Manager. Blackouts are a scheduled feature that instructs Grid Manager not to perform Discovery operations and Port Control/provisioning operations on the managed network at specified times, days and dates. See the section [Defining Blackout Periods](#) for details.

## Scheduling Additions and Modifications

You can schedule the addition and modification of an object. For example, you can schedule the addition of a DNS forward zone or the modification of a fixed address. After you schedule a task, administrators cannot modify the object associated with the scheduled task until after the appliance executes the task. However, the object can still be updated

with DHCP leases and other system services.

To schedule an addition or a modification:

1. Add or modify a record following the instructions described in this guide.
2. Click the Schedule icon at the top of the corresponding wizard or editor.
3. In the *Schedule Change* panel, complete the following:
  - **Now:** Select this to have the appliance perform the task when you save the entry. This is selected by default when there is no scheduled task associated with the object.
  - **Later:** Select this to schedule the task for a later date and time. Complete the following:
    - **Date:** Enter a date in YYYY-MM-DD (year-month-day) format. The appliance displays today's date. You can also click the calendar icon to select a date from the calendar widget.
    - **Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
    - **Time Zone:** Select a time zone for the scheduled date and time from the drop-down list. This field displays the time zone of the browser that the admin uses to log in to Grid Manager.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Scheduling Appliance Operations

The appliance supports the scheduling of the following operations:

- IP discoveries—For information, see [About Discovery](#).
- Service restarts—For information, see [Restarting Services](#).

## Scheduling Deletions

You can schedule the deletion of an object or an operation for a later date and time. However, you cannot schedule the deletion of a previously scheduled task.

To schedule a deletion:

1. Navigate to the object.
2. Select **Schedule Deletion** from the Delete drop-down menu.
3. In the *Schedule Deletion* dialog box, complete the following:
  - **Delete Now:** Select this to delete the object upon clicking **Delete Now**.
  - **Delete Later:** Select this to schedule the deletion at a later date and time. Complete the following:
    - **Date:** Enter the date in YYYY-MM-DD (year-month-day) format. The appliance displays today's date. You can also click the calendar icon to select a date from the calendar widget.
    - **Time:** Enter the time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
    - **Time Zone:** Select a time zone for the scheduled date and time from the drop-down list. This field displays the time zone of the browser that the admin uses to log in to Grid Manager.
4. Click **Schedule Deletion**.

The appliance performs the deletion at the scheduled date and time.

## Scheduling Recursive Deletions of Network Containers and Zones

Superusers can determine which group of users are allowed to schedule the deletion of a network container and its child objects as well as a zone and its child objects. For information about how to configure the recursive deletion of network containers and zones, see [Managing a Grid](#).

To schedule the recursive deletion of network containers and zones:

1. Navigate to the object.
2. Select **Schedule Deletion** from the Delete drop-down menu.
3. In the *Schedule Deletion* dialog box, complete the following:
  - **Delete Now:** Select this to delete the object upon clicking **Delete Now**.
  - **Delete Later:** Select this to schedule the deletion at a later date and time. Complete the following:

- **Date:** Enter the date in YYYY-MM-DD (year-month-day) format. The appliance displays today's date. You can also click the calendar icon to select a date from the calendar widget.
- **Time:** Enter the time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
- **Time Zone:** Select a time zone for the scheduled date and time from the drop-down list. This field displays the time zone of the browser that the admin uses to log in to Grid Manager. Select one of the following:
  - **Delete only the parent container:** Select this to delete only the parent objects and re-parent the child objects.
  - **Delete the parent container and its children:** Select this to delete the parent objects and all its child objects.

4. Click **Schedule Deletion**.

The appliance performs the deletion at the scheduled date and time.

## Viewing Scheduled Tasks

After you schedule a task, you can view the pending task in the **Administration** tab -> **Workflow** tab -> **Task Manager** tab. For more information, see [Viewing Tasks](#). Superusers can view all scheduled tasks, and limited-access admins can view their own scheduled tasks.

In certain panels such as the Network list panel and Smart Folders, Grid Manager displays a calendar icon next to objects that are associated with scheduled tasks, except for the addition of an object. You can click the icon to view the configuration and schedule. You can also reschedule the task if you are the owner of the task, a superuser, or an approver of the task (after you have approved it). In the corresponding editor, the Schedule icon is green when there is a pending scheduled task.

## Icons for Scheduled Tasks

Grid Manager displays a scheduled task icon next to an object that is associated with a scheduled task (except for the addition of an object), as shown in the below Icon for a Scheduled Task figure. When you mouse over the icon, an informational dialog box appears displaying the type of action, the date and time of the scheduled task, and the person who scheduled the task.

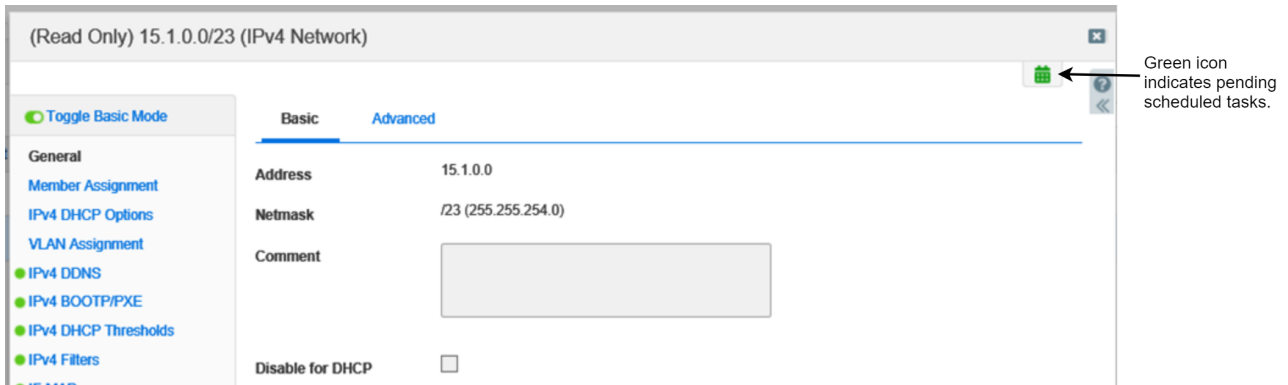
You can click the icon and Grid Manager displays the corresponding editor (for modification) or the *Scheduled Deletion* dialog box (for deletion) in the read-only mode. If you are viewing a task that you scheduled, you can modify and save the schedule, but you cannot modify the configuration of the object. If you are not the owner of a scheduled modification or a superuser, you can only view the information. You cannot reschedule the task. If you are not the owner of a scheduled deletion or a superuser, Grid Manager does not display the *Scheduled Deletion* dialog box when you click the icon.

*Icon for a Scheduled Task*



In the editor, Grid Manager displays the Schedule icon in green to indicate a pending scheduled task associated with the corresponding object, as shown in the Scheduling Icon Indicating a Pending Task figure. You can click the Schedule icon to view the date and time of the scheduled task. You can also reschedule the task if you have the applicable permissions. For information, see [Rescheduling Tasks](#).

*Scheduling Icon Indicating a Pending Task*



## Pending Tasks for Operations

You can view all pending tasks for a network discovery or service restart in **Task Manager** if you have the applicable permissions. For information, see [Viewing Tasks](#). You can also view the pending tasks in their corresponding dialogs. To view the pending tasks in an editor:

1. Network **Discovery**: From the **Data Management** tab, select the **IPAM** tab, and then click **Discovery** from the Toolbar.

**Service Restarts for the Grid**: From the **Data Management** tab, select the **IPAM**, **DHCP** or **DNS** tab, and then click **Restart Services** from the Toolbar, or from the **Grid** tab, click **Restart Services** from the Toolbar.

**Service Restarts for Grid members**: From the **Data Management** tab, select the **DHCP** or **DNS** tab -> **Members** tab, select a member checkbox, and then click **Restart Services** from the Toolbar.

2. Click the Schedule icon at the top of the wizard, and then select **Click here to view/manage the scheduled items**. Note that this link appears only when you have one or more scheduled tasks.

3. Grid Manager displays the following information in the *Scheduled Tasks*:

- **Scheduled Time**: The date, time, and time zone when the appliance executes the task.
- **Submitted Time**: The date, time, and time zone when the task was submitted.
- **Submitter**: The admin who scheduled the task.
- **Task Details**: The message that appears in the audit log.

By default, the appliance sorts the tasks by **Scheduled Time** starting with the earliest scheduled start time. You can do the following in this viewer:

- Sort the tasks in ascending or descending order by column, except for **Task Details**.
- Reschedule a selected task. For information, see [Rescheduling Tasks Associated with Operations](#).
- Delete a selected task by selecting the task checkbox and clicking the Delete icon.
- Export and print the information in the table.

## Rescheduling Tasks

Superusers can reschedule any scheduled task. Limited-access admins can reschedule only the tasks that they scheduled, depending on their permissions. Approvers can reschedule tasks that they have approved, if they have the scheduling permission. You can reschedule a task from different panels of Grid Manager, depending on your permissions. When you reschedule a task from the object list panel, Grid Manager displays the object or operation configuration in a read-only mode. You can modify the date and time to reschedule the task. However, you cannot modify the configuration of the object or operation. You can also reschedule your own task or a task you have approved from Task Manager.

To reschedule tasks associated with objects, see [ReschedulingTasksAssociatedWithObjects](#) below.

To reschedule tasks associated with operations, see [ReschedulingTasksAssociatedwithOperations](#) below.

## Rescheduling Tasks Associated With Objects

You can reschedule a task associated with an object from the *Scheduled Tasks* viewer or in an editor if you have the applicable permissions.

To reschedule a task from **Task Manager**:

1. From the **Administration** tab, select the **Workflow** tab -> **Task Manager** tab -> *scheduled\_task* checkbox, and then click the Reschedule icon.
2. In the *Reschedule* dialog box, modify the date and time when you want the appliance to execute the task. You can select **Now** to execute the task when you save the entry.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

To reschedule a task in an editor:

1. Navigate to the object with a scheduled task that you want to reschedule.
2. Click the scheduled task icon next to the object.
3. **For modification:** In the editor, click the Schedule icon at the top of the editor. In the *Schedule Change* panel, modify the date, time, and time zone. You can also select **Now** to execute the task upon saving the entry.  
**For deletion:** In the *ScheduleDeletion* dialog box, modify the date, time, and time zone. You can also select **Delete Now** to delete the object upon clicking **Delete Now**. The appliance puts the deleted object in the Recycle Bin, if enabled.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Rescheduling object change tasks and associated port control tasks is a special case allowing for rescheduling of both tasks in the same editor:

## Rescheduling an object change task and a port control task

auto\_arp\_refresh\_before\_switch\_port\_polling to=true cidr\_collection to=true finger\_printing to=true ipv4\_ping\_sweep to=true netbios\_scanning to=true port\_scanning to=true smart\_subnet\_ping\_sweep to=true snmp\_collection to=true switch\_port\_data\_collection\_polling to=PERIODIC switch\_port\_data\_collection\_polling\_interval to=3600 cidr: to=32 comment: common\_properties: domain\_name\_servers routers disabled: to=false discovery\_member: Member discovery.com enable\_discovery: to=true enable\_immediate\_discovery: to=true network\_view: NetworkView default use\_basic\_polling\_settings: to=false use\_member\_enable\_discovery: to=true vians:

Scheduled Time	Submitter	Affected Object	Object Type	Action
2020-04-30 15:01:14 IST	admin	EX2200-48T-4G	Device	Network Provisioning
Task Detail				
network_view: to=default interface: to=ge-0/0/0.0 router_ip: to=10.35.173.7 address: to=10.35.173.7 device: to=EX2200-48T-4G cidr: to=32 enable_dhcp_relay: to=false				

### Reschedule Object Change Task

Now  
 Later

**Selected time:**

**Date**

**Time**

**Time Zone**

**Your time:** 2020-04-29 03:06:48 PM  
 (UTC + 5:30) Bombay, Calcutta, Madras, New Delhi

### Reschedule Port Control Task

At same time as above  
 Later

**Selected time:**

**Date**

**Time**

**Time Zone**

**Your time:** 2020-04-29

Cancel
Save

## Rescheduling Tasks Associated with Operations

To reschedule a network discovery or a service restart:

1. From the **Administration** tab, select the **Workflow** tab -> **Task Manager** tab -> *scheduled\_task* checkbox, and then click the Reschedule icon.  
or  
Navigate to the operation and click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, select **Click here to view/manage the scheduled items**. Grid Manager displays all scheduled tasks related to the operation in the *Scheduled Tasks* viewer. Select the task checkbox, and then click the Reschedule icon.
2. Grid Manager displays detailed information about the task in the *Reschedule* dialog box.
3. Modify the date and time when you want the appliance to execute the task. You can also select **Now** to execute the task when you save the entry.
4. Save the configuration and click **Restart** if it appears at the top of the screen.



## Canceling Scheduled Tasks

To cancel a scheduled task:

1. From the **Administration** tab, select the **Workflow** tab -> **Task Manager** tab -> *scheduled\_task* checkbox, and then click the Delete icon.
2. In the *Confirm Delete Request* dialog box, click **Yes**.

The appliance deletes the scheduled task and does not perform the scheduled operation. Therefore, no change is made to any record after you delete a scheduled task.

## Performing Basic Operations Using Grid Manager

This section describes how to use some of the key NIOS features. It includes the following topics:

- [Logging on to the NIOS UI](#)
- [Setting Login Options](#)
- [Browser Limitations](#)
- [Multilingual Support](#)
- [Managing Internationalized Domain Names](#)
- [Finding and Restoring Data](#)
- [Opening Technical Support Requests](#)

### Logging on to the NIOS UI

The NIOS UI (also called Grid Manager) is the web interface that provides access to your appliance for network and IP address management. It provides a number of tools that you can use to effectively manage your appliance and IP address space.

You can log in to Grid Manager as long as you have permission to log in to the NIOS appliance. Superusers have unrestricted access to Grid Manager. Limited-access users require read-only or read-write permission to the data that they want to manage through Grid Manager. Grid Manager allows limited-access users to view and manage only the data for which they have permission. For example, to view IPv4 networks, you must have at least read-only permission to IPv4 networks, or to run a discovery, you must have read/write permission to the Network Discovery feature.

Note that superusers must configure admin groups and accounts in Grid Manager application of the NIOS appliance. In Grid Manager, superusers can set and change permissions for specific objects, such as IPv4 networks, IPv6 networks, and resource records. For information about user accounts and administrative permissions, see [Managing Administrators](#).

Before you log in to Grid Manager, ensure that you have installed your NIOS appliance as described in the installation guide, or the user guide that was shipped with your products and then configure your NIOS appliance accordingly. You must upload the CA certificate(s) that issued the client certificate to ensure a successful SSL/TLS connection to the appliance.

To log in to Grid Manager, complete the following:

1. Open an Internet browser window and enter **https://<IPaddress\_or\_hostname\_of\_your\_NIOSappliance>**. Grid Manager login page appears. For information, see [Supported Browsers](#).
2. Enter your user name and password, and then click **Login** or press the **ENTER** key. The default user name is **admin** and the default password is **infoblox**. Note that if your password expired or was reset by a superuser, you may be required to enter a new password.
  - If you are a smart card user and two-factor authentication is enabled on the appliance, your user name, which is the same as your CN (common name) in the client certificate, appears automatically. Enter the password you use to log in to the user account. For information about two-factor authentication, see [Authenticating Admins Using Two-Factor Authentication](#).

- In NIOS 8.5.2 or later, for a Grid Master or a standalone vNIOS instance deployed on AWS, you are prompted to reset the password on the first login attempt. You must reset the default password as a security requirement.
  - To reset, enter a new password and confirm it. The minimum password length must be four characters. It must consist of at least one uppercase character, one lowercase character, one numeric character, and one symbol character. Example: Infoblox1!.  
If the symbol character is at the beginning of the password, then include the password within quotes ("). Example: '@Infoblox123'.
  - Alternatively, you can define a new password in the **User data** field of the AWS console, in which case, you are not prompted to reset the password on the first login in Grid Manager. For more information, see *Provisioning vNIOS for AWS Using the BYOL Model* in the [Installation Guide for vNIOS for AWS](#) documentation.
- 3. Read the *Infoblox End-User License Agreement* (EULA), and then click **I Accept**.  
Note that if you want to view the privacy policy of Infoblox, then on the EULA screen, click **Infoblox Privacy Policy**. Grid Manager displays the policy on a new browser tab.
- 4. Click **OK**. The *Grid Setup* wizard appears.

## Setting Login Options

Grid Manager provides several options that you can set to facilitate the login process. Additionally, you can manage CA (Certificate Authority) and server certificates on the NIOS appliance. You can import certificates, select and view their details, or remove them. To manage certificates, see [Managing Certificates](#).

## Specifying Grid Name and Host Name

To define the default host name that appears when the login prompt displays, complete the following:

1. On the **Grid** tab -> **Grid Manager** tab, expand the Toolbar and select **Grid Properties** -> **Set up (Grid Setup Wizard)**.
2. On the Welcome page, select **Configure a Grid Master**, and then click **Next**.
3. Enter the Grid name in the **Grid Name** field and the host name in the **Host Name** field.

## Creating a Login Banner

You can create a statement that appears at the top of the *Login* screen (a banner message). This function is useful for posting security warnings or user-friendly information well above the username and password fields on the *Login* screen. A login banner message can be up to 3000 characters long. In a Grid, perform this task on the Grid Master.

To create a login banner, complete the following:

1. On the **Grid** tab -> **Grid Manager** tab, expand the Toolbar and select **Grid Properties** -> **Edit**.
2. In the *Grid Properties Editor*, on the **Security** tab, select **Enable Login Banner**. In the text field, enter the text that you want to be displayed on the login screen.
3. Save the configuration.

## Changing the Password and Email Address

Grid Manager creates and stores a user profile for each admin user. Each user profile contains information about the admin group and admin type assigned to the user. You can modify certain information in your user profile any time after the initial login. You can change your password to facilitate future logins and add your email address for reference.

Note that when multiple users log in to Grid Manager using the same admin account, they share the same user profile and preference settings, such as the widget, table size and column settings, independent of their browser settings. Instead of using the same admin account for multiple users, you can add multiple users to the same admin group so they can share the same permissions. For information about configuring admin accounts and admin groups, see [Managing](#)

## Administrators.

If you can access only the Tasks Dashboard, you may not see or configure certain fields in the *User Profile* editor.

To change your password and email address, complete the following:

1. At the top-right corner of the navigation bar, click the admin name and select **Profile** from the drop-down menu.
2. In the *User Profile* editor, complete the following:
  - **Name:** Displays your user name.
  - **Last Login:** Displays the timestamp of your last log in.
  - **Type:** Displays your user type. There are two user types: **Local** and **Remote**. Local type admin accounts are stored in the NIOS database. Remote type admin accounts are stored on another server, such as a RADIUS server. Grid Manager automatically deletes remote user profiles if users have not logged in for more than six months.
  - **Group:** Displays the admin group to which your account belongs. The admin group determines your administrative permissions. Only superusers can define admin groups through Grid Manager.
  - **Password:** You can set a new password according to the requirements that are displayed.
    - **Set Password:** If you are a local user, select this checkbox to set a new password for your account. If you are a remote user, this field does not appear.
    - **Old Password:** Enter your current password.
    - **New Password:** Enter the new password, and then re-enter it in the **Retype Password** field.
  - **Email Address:** Enter your email address. Note that this address simply provides contact information. By default, this field is blank.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Specifying the Table Size

You can specify the amount of data Grid Manager must display in a table or a single list view. You can improve the display performance by setting a smaller table size. The setting you specify here applies to all tables in Grid Manager. Note that if you can access only the Tasks Dashboard, you cannot configure table size.

To specify table size, complete the following:

1. At the top-right corner of the navigation bar, click the admin name and select **Profile** from the drop-down menu.
2. In the *User Profile* editor, in the **Table Size** field, specify the number of lines of data you want a table or a single list view to contain. You can set the number of lines from 10 to 256. The default is 20.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Selecting Your Home Page

When you first log in to Grid Manager, the Tasks Dashboard is your home page. You can change your home page for subsequent logins. You can specify the maximum number of widgets that can be configured per dashboard. You can set up to 20 widgets per dashboard. You can also set the auto refresh rate for dashboard widgets. This interval specifies how often the content of the dashboard widgets is refreshed.

To change your home page, complete the following:

1. At the top-right corner of the navigation bar, click the admin name and select **Profile** from the drop-down menu.
2. In the *User Profile* editor, complete the following:
  - **Default Dashboard:** Select **Status** or **Task** from the drop-down list.
  - **Maximum Widgets per Dashboard:** Specify the maximum number of widgets that can be configured per dashboard. You can enter a value between 1 and 20. The default value is 10. This limit does not apply to the default dashboard.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Grid Manager displays the selected dashboard as your home page when you log in the next time.

## Setting the Browser Time Zone

You can specify the time zone that Grid Manager uses to convert all displayed time values, such as the last discovered and last login time. Grid Manager sets the time zone based on the time zone of your browser when you set the time zone to auto-detect in the *User Profile* editor. When you set the time zone of your browser to auto-detect and Grid Manager cannot automatically determine the time zone when you log in, the time zone is set to the UTC (Coordinated Universal Time) standard. In this case, you can manually change the time zone in the *User Profile* editor.

To manually set the time zone of your browser, complete the following:

1. At the top-right corner of the navigation bar, click the admin name and select **Profile** from the drop-down menu. The *User Profile* editor displays your username, user type, and admin group.
2. In the *User Profile* editor, select the time zone Grid Manager uses to convert all displayed time values. The default is **Auto-detect time zone**. You must select a specific time zone when Grid Manager cannot automatically detect the time zone of your browser.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Browser Limitations

- When you use Internet Explorer 7 or 8 without installing the latest updates, Grid Manager may stop loading a page when you navigate from one tab to another or when you use the back navigation button to go back to a previous page. To solve this problem, you can press Ctrl+F5 to refresh the browser or install the latest updates.
- When you use the zoom function in Internet Explorer 7 running on Microsoft Windows XP, Grid Manager may not properly display some pop up windows. This is a known issue in Internet Explorer 7.
- In Internet Explorer 8, Grid Manager does not display the directory path of an uploaded file. Instead, it displays "fakepath" in place of the directory path. To resolve this issue, you can add Grid Manager as a trusted site or enable the "Include local directory path when uploading files to a server" feature in the browser. For information, refer to the MSDN documentation at <https://msdn.microsoft.com/en-us/library/ms535128.aspx>.
- When you use FireFox to access Grid Manager, tooltips do not display for disabled drop-down menu items. In addition, when you run a large query of smart folders, Grid Manager may display a warning message about "Unresponsive Script". Click **Continue** to proceed.
- Depending on the browser you use, Grid Manager may display a dialog box that indicates the system is unavailable during a system restart or reboot.
- Infoblox strongly recommends that you do not log in to Grid Manager from different browser windows using the same user account. Depending on the browser you use, it may cache user information in one session and apply it to another session. This can cause inconsistent behaviors within the browser sessions.

## Multilingual Support

The NIOS appliance supports UTF-8 (Unicode Transformation Format-8) encoding for the following:

- Hostnames for Microsoft Windows clients that support Microsoft Windows code pages. For information about Configuring UTF-8 Encoding for Hostnames, see [Configuring UTF-8 Encoding for Hostnames](#).
- Input fields through Grid Manager.

UTF-8 is a variable-length character encoding standard for Unicode characters. Unicode is a code table that lists the numerous scripts used by all possible characters in all possible languages. It also has a large number of technical symbols and special characters used in publishing. UTF-8 encodes each Unicode character as a variable number of one to four octets (8-bit bytes), where the number of octets depends on the integer value assigned to the Unicode character. For information about UTF-8 encoding, refer to [RFC 3629 \(UTF-8, a transformation format of ISO 10646\)](#) and the *ISO/IEC 10646-1:2000 Annex D*. For information about Unicode, refer to *The Unicode Standard*. Depending on the OS (operating system) your management system uses, you must install the appropriate language files in order to enter information in a specific language. For information about how to install language files, refer to the documentation that comes with your management system.

## UTF-8 Supported Fields

The NIOS appliance supports UTF-8 encoding in all of the comment fields and most input fields. You can enter non-English characters in these data fields through Grid Manager and the Infoblox API. When you use the Infoblox API, all the non-ASCII strings must be UTF-8 encoded so that you can use Unicode characters. The NIOS appliance does not support UTF-8 encoding for data that is configurable through the Infoblox CLI commands.

In general, the following items support UTF-8 encoding:

- All the predefined and user-defined extensible attributes.
- All the comment fields in Grid Manager.
- File name fields for FTP and TFTP backup and restore operations.
- The login banner text field. When you use the serial console or SSH, the appliance cannot correctly display the UTF-8 encoded information that you enter for the login banner.



### Note

For data fields that do not support UTF-8 encoding, the appliance displays an error message when you use non-English characters.

## UTF-8 Support Limitations

The NIOS appliance has the following UTF-8 support limitations:

- Object names that have data restrictions due to their usage outside of the Infoblox database do not support UTF-8 encoding. For example, IP addresses and Active Directory domain names.
- When importing a database, most of the ASCII control characters cannot be encoded. This might cause failures in upgrades or database restore operations.
- Search is based on the Unicode standard. Depending on the language, you might not be able to perform a case-sensitive search.
- Binary data is encoded as text.
- UTF-8 encoding does not fully support regular expressions. It matches constant strings. However, it does not encode characters that are inside square brackets or followed by regular expressions such as \*, ?, or +.
- You can use UTF-8 characters to authenticate both the User Name and Password through the Infoblox GUI, but not through the Infoblox CLI.
- You cannot use UTF-8 characters to name a custom DHCP fingerprint. For information about DHCP fingerprint detection, see [DHCP Fingerprint Detection](#).

## Managing Internationalized Domain Names

The Infoblox Grid supports IDNs (Internationalized Domain Names) for DNS zones and resource records to provide the flexibility of specifying domain names in non-English characters.

An IDN is a domain name that contains a language-specific script or alphabet, such as Arabic, Chinese, Russian, Devanagari, or the Latin alphabet-based characters with diacritics, such as French. IDNs are encoded in multi-byte Unicode and are decoded into ASCII strings using a standardized mechanism known as Punycode transcription. For example, DNS Zone 'инфоблокс.рф' (IDN in Russian) can be written as 'xn--90anhdigczv.xn--p1ai' in the punycode representation. In addition, the appliance has a built-in conversion tool to assist you in identifying and troubleshooting an IDN or the punycode representation of an IDN. For information about how to decode IDNs, see [Decoding IDNs and Encoding Punycode](#) below. The appliance supports IDNs in certain fields. There are certain guidelines and limitations about IDN support.

## Decoding IDNs and Encoding Punycode

You can encode non-English characters into punycode and decode punycode to obtain a domain name in its original character set. You can encode IDNs and decode punycode simultaneously. You can use special characters.

To encode non-English character set into punycode and decode punycode:

1. Select any tab in Grid Manager, and then click **IDN Converter** from the Toolbar.
2. In the *IDN Converter* wizard, complete the following:
  - Specify the domain name in the **Unicode** text box and click **Convert to Punycode**. The **Punycode** field displays the punycode representation of the domain name.
  - Specify the punycode representation of a domain name in the **Punycode** field and click **Convert to Unicode**.  
The **Unicode** field displays the domain name in its original character set.  
Note that you can use special characters in the **Unicode** and **Punycode** fields.
  - Click **Clear** to clear the entries. Note that when you click **Clear** for a specific conversion, the appliance clears only the error message that corresponds to that conversion.
3. Click **Close**.

## IDN Supported Fields

The NIOS appliance supports IDNs in all domain name fields. For information about IDNSupportForDNS Zones, see [Domains and Zones](#). You can enter non-English characters in the domain name fields through Grid Manager and the Infoblox API. The NIOS appliance does not support IDNs for data that is configurable through the Infoblox CLI commands. You can use the punycode representation to configure data through the CLI commands.

The appliance supports IDNs in the following:

- You can use UTF-8 characters when defining your own hostname checking policy. For information, see [Specifying Hostname Policies](#).
- You can use both IDNs and punycode to search for IDN data through Global Search. For more information about using Global Search, see [Finding and Restoring Data](#).
- Use smart folders to organize and monitor IDN data. However, if the content in a smart folder contains IDNs, then the punycode representation is not available. For information, see [Smart Folders](#).
- You can import data that contains IDNs in CSV format for the supported fields and objects using CSV import. For more information, see [Importing and Exporting Data using CSV Import](#). For a list of supported record types and specific guidelines for creating a data file, refer to the *Infoblox CSV Import Reference*.
- The IPAM tab displays IDNs for DNS resource records associated with IP addresses, such as A records, AAAA records, hosts, and PTR records. For information, see [About IP Address Management](#).
- The audit log entries are displayed in their original characters. The audit log contains IDN data as received by the appliance and as specified by the administrators. Note that the punycode representation generated by NIOS is not displayed in the audit log.
- When you upgrade from a previous NIOS release, the appliance converts all punycode to IDNs. If the conversion fails, the appliance retains the punycode representation to avoid upgrade failure. For information about upgrades, see [Guidelines for Upgrading, Backing Up, and Restoring Data](#).
- When you restore a backup file from a previous NIOS release, the appliance converts all punycode to IDNs. If the conversion fails, the appliance retains the punycode representation to avoid failure to restore the database. For information, see [Guidelines for Upgrading, Backing Up, and Restoring Data](#).
- If synchronized data between the appliance and Microsoft servers contains IDNs, the IDNs are preserved. For information, see [Managing Microsoft DNS Services](#).

## IDN Support Limitations

The appliance has the following IDN support limitations:

- F5® load balancers does not support IDNs. The NIOS appliance does not encode punycode to IDNs for F5 load balancer related objects. Only the punycode representation is available.
- Multi-Grid configuration does not support IDNs.

- The Infoblox CLI does not support IDNs.
- If a resource record containing an IDN is added to the Infoblox Grid through DDNS updates, the domain name field displays the record name in UTF-8 encoded format. For more information, see [Managing Resource Records](#).
- The following FQDNs does not support IDNs:
  - FQDN of an external DNS Server (direct or via name server group)
  - FQDN of a DNS root server
  - FQDN of a Microsoft server
  - FQDN of an Infoblox Grid Member
  - FQDN of an external authentication source (Active Directory, LDAP, OCSP, RADIUS, TACACS+)
  - FQDN of an NTP server
  - FQDN of a Thales Luna HSM Module
  - FQDN of an email relay server
  - FQDN of a vSphere/ESX server
  - FQDN of a Kerberos Key Distribution Center

## Using IDNs for Unsupported Objects

The appliance accepts only punycode entries for objects that do not support IDNs. To use IDNs for these objects, manually convert IDNs to punycode and use the punycode representation.

Use the punycode representation of IDNs for the following:

- When you configure domain names in forwarder servers, NXDOMAIN rulesets, blacklist rules, and DNS resolver search lists.
- When you configure domain names for DHCP and DHCPv6 services, including DDNS domain name, any DHCP options that accept domain names (host-name (12) string) or lists of domain names (domain-search (119) domain-list), and DHCPv6 options that accept domain names (dhcp6.fqdn (39) string) or lists of domain names (dhcp6.domain-search (24)) domain-list.
- When you add domains in the Inclusion list and Exclusion list. For information about Excluding Domains From Query and Response Capture, see [Capturing DNS Queries and Responses](#).
- When you configure rules for a local RPZ and RPZ feed. For information, see [Configuring Local RPZs](#) and [Configuring Infoblox Threat Intelligence Feed](#).

## Displaying IDN Entries in Punycode

The appliance displays IDN entries in punycode for the following:

- The data of a zone for which an Infoblox Grid member is the secondary server.
- The CLI commands `dig`, `ddns_add`, `ddns_delete`, `show dns`, and `set dns` support punycode only. For information about CLI commands, refer to the [Infoblox CLI Guide](#).
- All syslog entries generated by DNS.
- IDN data in database files is stored in punycode.
- The DNS cache of a Grid member that contains IDNs.
- The **Reporting** tab displays all report data that contains IDNs in punycode. For information, see [Predefined Dashboards](#).

## Finding and Restoring Data

Grid Manager provides tools for organizing and quickly retrieving your DNS, DHCP and IP address management data.

The *Finder* panel, which appears on all pages in Grid Manager, provides tools for organizing your data. The *Finder* panel provides easy access to the following:

- [Using Bookmarks](#)
- [Using the Recycle Bin](#)
- [Managing Third Party URL Links](#)
- [Using Filters](#)
- [Using Quick Filters](#)



- [Using Global Search](#)
- [Using the Go To Function](#)
- Smart Folders: Contains a hierarchical list of smart folders that are available in My Smart Folders. For information, see [Smart Folders](#).
- Bookmarks: Contains bookmarked objects, such as networks and IP addresses.
- Recycle Bin: Contains deleted objects that can be restored or permanently removed.
- URL Links: Contains a list of third-party URLs that you previously added. You can add more URL links, and modify and delete existing URL links. For information, see [Managing Third Party URL Links](#) below.

In the *Finder* panel, you can expand and collapse these sections. To expand a section, click the + icon next to the header. To collapse a section, click the - icon.

In addition, Grid Manager also provides the following:

- Filters to customize data displays. For more information, see [Using Filters](#) and [Using Quick Filters](#) below.
- Global Search to search the NIOS database for objects that match your criteria. For more information, see [Applying Quick Filters](#) below.
- Go To function to quickly locate an object. For more information, see [Using the Go To Function](#) below.

## Using Bookmarks

The Bookmarks section displays objects for which you have created bookmarks. You can create bookmarks for objects such as networks, DNS zones, and admin groups. To bookmark an object, navigate to its page and click the Bookmark icon at the top of the page. If you have more than one network view, Grid Manager displays the name of the bookmark with the network view to which the object belongs. For example, when you bookmark IP address 10.128.0.10 in the default network view, Grid Manager displays the bookmark as default > 12.128.0.10.

However, if you have only one network view, Grid Manager displays only the object name 12.128.0.10. If you create a bookmark before adding more network views, the bookmark name (without the network view) remains the same. You can rename the bookmark at any time. You can create only one bookmark for each object, up to 500 objects. When your bookmarks are close to 500, you may want to remove some to create room for new ones.

You can perform the following in Bookmarks:

- Access a bookmarked object
- Edit the name of a bookmark
- Delete a bookmark

To access a bookmarked object, click the name of the bookmark. Grid Manager displays the network view to which the bookmarked object belongs. For example, clicking on the bookmark of network 10.0.1/24 takes you to the network list view. You cannot access an object that has been deleted.

You can arrange the order of the bookmarked objects by dragging and dropping the objects in the *Finder* panel.

To edit the name of a bookmark:

1. Mouse over to the bookmark.
2. Click the Edit icon.
3. Modify the name of the bookmark. Note that you cannot create multiple bookmarks with the same name.

To delete a bookmark:

1. Mouse over to the bookmark.
2. Click the Delete icon. Grid Manager removes the bookmark.

## Using the Recycle Bin

The Recycle Bin section contains objects that you deleted. It provides a way to restore data where the deletion of the object (such as a network) could result in a major data loss.

You must enable the Recycle Bin in Grid Manager to store and restore deleted objects. For information about how to enable and disable the Recycle Bin, see [Enabling and Disabling the Recycle Bin](#) below. When you use the Recycle Bin, you can restore deleted objects to the active configuration. You can also permanently remove the objects from the



Recycle Bin. If you do not enable the Recycle Bin, the appliance immediately removes objects from the database when you delete them using Grid Manager.



#### Note

When you upgrade to a new NIOS release, the appliance permanently deletes the objects from the Recycle Bin.

On a NIOS appliance, only superusers have permissions to fully manage the Recycle Bin. If you have limited-access permissions, you can view, restore, and permanently remove only the objects that you deleted. For Cloud Network Automation, the Recycle Bin is not supported on the Cloud Platform Appliance. Only deletions performed on the Grid Master are stored in the Recycle Bin. Deleted objects can only be restored from the Grid Master. For information about Cloud Network Automation, see [Deploying Cloud Network Automation](#).

You can do the following in the Recycle Bin:

- View deleted objects
- Restore deleted objects
- Remove deleted objects
- Empty the Recycle Bin

### Enabling and Disabling the Recycle Bin

To enable or disable the Recycle Bin:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties* Editor, select the General tab, and then complete the following:  
Select **Enable Recycle Bin** to enable the Recycle Bin  
or  
Deselect **Enable Recycle Bin** to disable the Recycle Bin.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

### Viewing Objects in the Recycle Bin

Grid Manager displays the short name of all deleted objects in the Recycle Bin. For example, the short names for hosts and resource records are their domain names, and the short names for fixed addresses and reservations are their IP addresses.

The Recycle Bin does not display all deleted objects; it can display up to 15 of the most recently deleted objects. When the Recycle Bin contains objects that are not displayed in the *Finder* panel or multiple objects that have the same name, the **Show All** button appears. Click the button to display the *Recycle Bin* dialog box that contains detailed information about each deleted object. When you have multiple deleted objects that use the same name, you may want to view detailed information about the deleted objects before taking any action. You can remove and restore selected objects and empty the Recycle Bin in the *Recycle Bin* dialog box.

You can do the following in the *Recycle Bin* dialog box:

- Sort the data in ascending or descending order by column.
- Use filters and the search function to look for specific objects. For information about filters, see [Using Filters](#) below.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#) below.
- Use the **GoTo** function to quickly find the data in the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **GoTo** field and select the object from the possible matches.

To view detailed information about deleted objects:

1. In the *Finder* panel, expand **Recycle Bin**.
2. Click **Show All**.

Grid Manager displays the *Recycle Bin* dialog box that contains the following information for each object:

- **Name:** The short name of the object. For example, the short names for fixed addresses and reservations are their IP addresses.

- **Type:** The object type.
- **Parent/Container:** The parent object or parent container to which the object belongs.
- **Admin:** The admin name of the user who deleted the object.
- **Data:** The data that the object contains, if any.
- **Network View:** The network view to which the object belongs.
- **Time:** The time stamp when the object was deleted. To close the dialog box, click **Close**.

## Restoring Objects from the Recycle Bin

You can restore deleted objects from the Recycle Bin only if you enable the Recycle Bin, and only if you select an object in the panel. You can restore only one object at a time. Deleted objects are stored in the Recycle Bin until you delete them or empty the bin.

To restore items from the Recycle Bin:

1. In the *Finder* panel, expand **Recycle Bin**.
2. Select the object you want to restore.
3. Click the Restore icon.

Grid Manager restores the object to its corresponding container or configuration. You can confirm the restoration by checking that the object does not appear in the Recycle Bin any longer, and that it is reestablished in the appropriate panel in the GUI.

## Deleting Objects in the Recycle Bin

You can permanently delete individual objects in the Recycle Bin only if the Recycle Bin is enabled.

To delete objects in the Recycle Bin:

1. In the *Finder* panel, expand **Recycle Bin**.
2. Select the object you want to delete.
3. Click the Delete icon.  
Grid Manager displays the *Confirm Delete* dialog box.
4. Click **Yes** to delete the object.

## Emptying the Recycle Bin

You can permanently delete the contents of the Recycle Bin, if enabled. Only superusers can empty the Recycle Bin. Because the Recycle Bin can grow large, you can periodically empty the Recycle Bin to free up disk space.

To empty the Recycle Bin:

1. In the *Finder* panel, expand **Recycle Bin**.
2. Click **Empty**.  
Grid Manager displays the *Confirm Empty Recycle Bin* dialog box to confirm that you wish to empty the Recycle Bin.
3. Click **Yes**.

## Managing Third Party URL Links

In the URL Links section, you can add the URL links of frequently used third party portals and destinations.

### Adding URL Links

To add URL links:

1. In the *Finder* panel, expand **URLLinks**.
2. Click **Add**.

3. In the *URL Configuration* dialog box, complete the following:
  - **URL:** Enter the URL of the destination characters. When you enter the URL, the appliance validates the entry. You cannot save the entry if the URL is not in a valid format.
  - **Name:** Enter a name that represents the portal or site of the URL.
  - **Set as global parameter:** This field appears only if you log in as a superuser. Select this checkbox to make the URL link globally available to all users.
  - **Logo:** Click **Upload** to add a logo to the URL. The appliance displays the logo in 16x16 pixels. Click **Reset to Default** to use the default logo.
4. Save the configuration.

## Modifying URL Links

To modify the information you entered for an existing URL link:

1. In the *Finder* panel, expand **URLLinks**.
2. Hover your mouse over the URL you want to modify, and then click the Edit icon.
3. In the *URL Configuration* dialog box, modify the information as described in [Adding URL Links](#).

## Deleting URL Links

To permanently delete an URL link:

1. In the *Finder* panel, expand **URLLinks**.
2. Hover your mouse over the URL you want to delete, and then click the Delete icon.
3. In the *Delete URL Link* dialog box, click **Yes**.

## Using Filters

You can control the amount and the kind of data displayed in a specific panel by adding filter criteria. When you add filter criteria, the appliance screens the data based on your filter rules and displays only the information that matches the rules. To narrow your search for specific information, you can add up to 10 filter rules. In some panels, such as the DHCP Networks tab, you can switch between viewing information with and without the filter criteria by toggling the filter on and off. You can save filter criteria as quick filters so you can reuse the same filter rules to obtain updated information without redefining them each time you log in to the appliance.

You can also use filters to find objects that have failed an operation. When you try to modify multiple objects with the same extensible attribute, the appliance may not modify all of the selected objects. For information about extensible attributes, see [Managing Extensible Attributes](#). For example, after you modify the extensible attribute "Building" with new value "West", you can find the objects that are not updated by defining a filter with "Building" "does not equal" "West".

Depending on the filter criteria, you can use different filter operations to narrow down your search results. Grid Manager supports the following filter operations based on your selected filter criteria:

- **equals:** Defines a specific value for a selected filter criterion
- **does not equal:** Defines a selected filter criterion that does not equal a specific value
- **begins with:** Specifies a beginning value for a selected filter criterion
- **does not begin with:** Specifies a selected filter criterion that does not begin with a specific value
- **has a value:** Specifies a selected filter criterion that contains a value
- **does not have a value:** Specifies a selected filter criterion that does not contain a value
- **belongs to:** Defines a selected filter criterion that belongs to a specific parent object
- **Inheritance State equals:** Specifies a specific inheritance state

To use a filter:

1. In a panel, click **Show Filter** to enable the function.
2. In the filter section, complete the following:
  - In the first drop-down list, select a field such as an object name, comment, or an extensible attribute (fields with a gray background) as the filter criterion. Grid Manager displays only the supported fields.
  - In the operator drop-down list, select an operator for the filter criterion. Depending on what you select in the first filter field, this list displays the relevant operators for the selection. The operator **Inheritance State**

- equals** is displayed only when you select an inheritable extensible attribute from the **Type** drop-down list. This operator is not displayed if the extensible attribute is not inheritable.
  - In the value field, enter or select the attribute value for the first filter field. Depending on what you select for the first two filter fields, you can either enter a value or select an attribute from a drop-down list. For example, if you select an extensible attribute in the first filter field, you can enter the attribute value here. If you select an inheritable extensible attribute from the **Type** drop-down list, and select **Inheritance State equals** in the operator drop-down list, the value field displays a drop-down list with these values: **Inherited** and **Overridden/No Parent**. When you select **Inherited**, extensible attributes that are inherited by the descendants are listed. When you select **Overridden/No Parent**, extensible attributes which are overridden or do not have a parent are listed.
- 3. Optionally, click the **+** icon to add another filter rule. You can add up to 10 filter rules.
- 4. Click **Apply** to apply the rules
  - or
  - Click **Reset** to clear the filter criteria.

To view information with or without the filter criteria:

- Click **ToggleFilterOn** to apply filter criteria to the displayed data. Grid Manager displays only the filtered data in the panel.
- or
- Click **ToggleFilterOff** to have the appliance list all data without applying filter criteria.

## Using Quick Filters

A quick filter saves filter rules that you define in a specific panel. You can reuse a quick filter to find updated information in a panel without specifying the same rules each time. Superusers can define quick filters and share them with local users. Limited-access users can only create quick filters for their own use. You can create up to 10 global and 10 local quick filters in each panel that supports filters.

The appliance supports the following quick filters:

- System quick filters: These are predefined filters. You cannot modify the criteria of these filters. System quick filters are prefixed with **[S]** in the quick filter list. Infoblox currently supports the following system quick filters in the DNS data panels:
  - All Forward Mapping Zones:** This quick filter displays all forward mapping zones in lexicographical order.
  - All Reverse Mapping Zones:** This quick filter displays all IPv4 and IPv6 reverse mapping zones in numerical order. The appliance displays IPv4 zones before IPv6 zones.
  - All IPv4 Reverse Mapping Zones:** This quick filter displays only the IPv4 reverse mapping zones in numerical order.
  - All IPv6 Reverse Mapping Zones:** This quick filter displays only the IPv6 reverse mapping zones in numerical order.
  - RPZ Logs:** This quick filter displays only the RPZ syslog messages in CEF format. This option is displayed only in the **Syslog** when RPZ license is enabled.
- Global quick filters: Only superusers can define global quick filters. You can make global filters available to all users. Limited-access users can use global quick filters, but they cannot modify them. Global filters are prefixed with **[G]** in the filter list.
- Local quick filters: Limited-access users can create local quick filters for their own use. You cannot share local quick filters with other users in the Grid. Local filters are prefixed with **[L]** in the filter list.



### Note

In the default DNS zone view, the appliance displays forward mapping zones first, followed by IPv4 reverse mapping zones, and then IPv6 reverse mapping zones.

## Adding Quick Filters

To add quick filters:

- In a panel that supports filters, click **Show Filters**.
- In the filter section, define filter criteria for the quick filter, as described in Using Filters.

3. Click **Save**.
4. In the *Save Quick Filter* dialog box, complete the following:
  - **Name:** Enter a name for the quick filter. The name must be 20 characters or longer. Ensure that you use a unique name for each quick filter in a particular filter category. For example, you can use the same filter name for both a global and local filter, but you cannot do so for two local filters.
  - **Set as a global quick filter:** This displays only if you log in as a superuser. Select this checkbox to make the quick filter globally available to all users.
5. Save the configuration.

The appliance adds the quick filter to the quick filter drop-down list in the specified panel.

## Modifying Quick Filters

To modify quick filters:

1. In a panel that supports filters, click **ShowFilters**, and then select the quick filter you want to modify from the **QuickFilter** drop-down list.
2. In the filter section, click the Edit icon next to the filter name.
3. Modify the filter criteria, as described in *Using Filters*.
4. Click **Save**.
5. In the *SaveQuickFilter* dialog box, you can click **Save** to save the modified filter criteria under the same quick filter name. You can also modify the quick filter name, as described in *Modifying Quick Filters*, and save the entry as a new quick filter.
6. Save the configuration.

## Applying Quick Filters

To apply quick filters:

1. In a panel that supports filters, click **Show Filters**, and select the quick filter from the **Quick Filter** drop-down list.
2. Based on the filter criteria, the appliance displays the filtered information in the panel. The selected quick filter remains active in the panel until you select another quick filter.

## Turning Off Quick Filters

You can do one of the following to turn off a quick filter:

- Select **None** from the quick filter drop-down list.
- Click **ToggleFilterOff** or **Reset** in the filter section.
- Delete a quick filter, as described in *Deleting Quick Filters*.

## Deleting Quick Filters

To delete quick filters:

1. In a panel that supports filters, click **ShowFilters**, and then select the quick filter you want to delete from the **QuickFilter** drop-down list.
2. In the filter section, click the Delete icon next to the filter name.
3. In the *DeleteQuickFilter* dialog box, click **Yes** to permanently delete the quick filter.

## Using Global Search

You can use the global search function to search the entire NIOS database for data that matches a specific value and filter criteria. You can define filter criteria and enter applicable search values to refine the search. Grid Manager supports regular expressions in global search. Grid Manager can display up to 500 search results. When search results exceed 500, a warning message appears and you may want to refine your search. Search results remain in the *Search* dialog box until you reset the search parameters or log out of Grid Manager. You can search for DNS zones and resource records that contain IDNs. For information about IDNs, see [Managing Internationalized Domain Names](#). Based on your search requirements, you can choose to perform a basic search or an advanced search.

A basic global search provides a faster way to locate frequently searched results using one specific filter criterion. It is designed to handle a large amount of data in an efficient manner. Supported filtering objects are DNS name (FQDN or CNAME only), DUID, IP address and MAC address. This is the recommended global search method if you have a large data set and only need to search by a single filter criterion. You cannot use regular expressions in a basic global search. An advanced global search allows you to perform complex searches by defining multiple filter criteria. You can add up to 10 filtering rules. You can also include existing extensible attributes for the matching objects. Note that if a search result contains duplicate records, the appliance displays only one record and discards others. For example, if the canonical name matches an alias name, the appliance displays only one CNAME record in the result.

 **Note**

- Depending on the size of your database, global search may take a long time to complete. Grid Manager times out when queries or searches take longer than 120 seconds. To expedite searches, use filters to refine the search criteria. You can also use basic global searches if you have a large data set and you only need to search by a single filter criterion.
- To search for any DHCP objects or its attributes using Global Search, you must have DHCP license installed.

You can also do the following in the Results table:

- Click the Open icon to view detailed information of the matching object.
- Click the Edit icon to edit the matching object information. For information about [Editing Matching Objects in Search Results](#), see below.
- Click the Extensible Attributes icon to edit the value of the respective extensible attribute. For information about [Editing Multiple Extensible Attributes in Search Results](#), see below.
- Click the Export icon to export the data displayed in the Results table.
- Click the Print icon to print the data displayed in the Results table.

To perform a global search:

1. Click the Global Search icon on the navigation bar.  
By default, the appliance opens the **Basic** tab of the *Search* dialog box.
2. Do one of the following:
  - If you want faster search results and you can search by DNS name, DUID, IP address, or MAC address, do the following in the **Basic** tab:

**Include Network Insight Devices and Interfaces:** This appears only when you have the Network Insight license installed. Select this checkbox to include devices and interfaces discovered through Network Insight. Note that it might take longer than expected for the appliance to return results for these objects. The setting is saved between user sessions.

    - **ChooseFilter:** Select a value from the drop-down list. You can explicitly search by DNS name (FQDN or CNAME only), DUID, IP address, or a MAC address. Note that you can apply only one filter at a time.
    - **ChooseOperator:** Select an operator for the filter criterion. Depending on what you select in the first filter field, this list displays the relevant operators for the selection. Possible values include **equals**, **beginswith**, and **contains**. For example, if you choose a **DNSName** filter, **contains** is the only applicable operator.
    - In the value field, enter the value for the filter field. For example, if you select DUID in the first filter field, you can enter the DUID of the client's device in the value field.
  - If you have a more complex search requirement and need to include multiple filter criteria, do the following in the **Advanced** tab:
    - In the first field, enter the value that you want your search results to match. For example, if you want to search for hostnames that contain "Infoblox," enter **Infoblox** in this field. You can also specify the value of an inheritable extensible attribute. You can use regular expressions in the search value.
    - In the **Type** drop-down list, select an object type, comment, or an extensible attribute (fields with a gray background) as the filter criterion. Grid Manager displays all the supported fields in the drop-down list. The default is **Type**. Grid Manager searches all objects when you use the default. You can narrow down the search and improve the search performance by selecting an object type. Extensible attributes are displayed with a gray background.

- In the operator drop-down list, select an operator for the filter criterion. Depending on what you select in the first filter field, this list displays the relevant operators for the selection. The operator **InheritanceStateequals** is displayed only when you select an inheritable extensible attribute from the **Type** drop-down list. This operator is not displayed if the extensible attribute is not inheritable.
  - In the value field, enter or select the attribute value for the first filter field. Depending on what you select for the first two filter fields, you can either enter a value or select an attribute from a drop-down list. For example, if you select an extensible attribute in the first filter field, you can enter the attribute value here. If you use the default **Type** in the first filter field, you can select an object or record type from the drop-down list. The default is **ALL**. Grid Manager searches all object types when you use the default. If you select an inheritable extensible attribute from the **Type** drop-down list, and select **InheritanceStateequals** in the operator drop-down list, the value field displays a drop-down list with these values: **Inherited** and **Overridden/NoParent**. When you select **Inherited**, extensible attributes that are inherited by the descendants are listed. When you select **Overridden/NoParent**, extensible attributes which are overridden or do not have a parent are listed.
  - Optionally, click the **+** icon to add another filter. You can add up to 10 filter rules.
  - **IncludeExtensibleAttributesValues**: Select this checkbox to include extensible attributes in the search results for the matching objects. Once selected, this configuration affects all future searches for the current user. Note that it might take longer for the search results to appear if there are a large number of extensible attributes associated with the matching objects.
3. Optionally, you can click **Reset** to clear the search results and start a new search. You can also click the Refresh icon to refresh the search results.  
Grid Manager stores the search results until you reset the search parameters or log out.
  4. After you finish defining filters, click **Search** or press **Enter**.



#### Note

You can save each search that contains multiple filter criteria as a quick filter for future use. For information about quick filters, see [Using Quick Filters](#) above.

In the Results table, Grid Manager displays the following information:

- **Name**: The name of the matching object. This field displays the name of the matching object and the path to the matching object if the object is a network or an IP address. You can click the link to open, view, and edit the object.
- **Type**: The type of the matching object. For example, bulk host, NS record, forward-mapping authoritative zone, or network container.
- **MatchedProperty**: The attribute or property of the matching object. For example, if the search value matches the email address that corresponds to a hostname, this field displays **Email**. If the search value matches the DNS view of a resource record in a DNS zone, this field displays **DNSView/FQDN**.
- **MatchedValue**: The value of the matching object. For example, if an IP address contains the search value, this field displays the IP address. If a hostname contains the search value, this field displays the hostname.
- **IPAddress**: The IP address of the matching object. When you click the IP address link, Grid Manager displays the corresponding IP address panel from which you can view detailed information.
- **Comment**: Comments that were entered for the matching object.
- **Site**: Values that were entered for the matching object.



#### Note

If you have selected to include extensible attribute values, you can select the corresponding columns to be displayed in the search results. Extensible attribute columns are hidden by default.

## Editing Matching Objects in Search Results

Grid Manager displays search results in the Results table. You can open and view detailed information about an object. You can also edit the properties of a selected object.

To edit an object in the Results table:

1. In the Results table, select the object checkbox.



2. Click the Open or Edit icon. You can also click the link of an object if Grid Manager displays the path. Grid Manager displays the object in the corresponding editor depending on the type of object you selected.
3. Edit the properties of the object in the editor.
4. Save your changes.

## Deleting Matching Objects in Search Results

You can delete one or multiple matching objects in the search Results table.

To delete a matching object:

1. In the Results table, select the object checkbox. You can delete multiple objects.
2. Click the Delete icon.
3. In the *Delete Confirmation* dialog box, click **Yes**.

Grid Manager deletes the selected objects from the database. Most deleted objects are stored in the Recycle Bin. You can print search results. You can also export search results in CSV (comma separated value) format. For information about CSV Import and Exporting Displayed Data, see [Importing and Exporting Data using CSV Import](#).

## Editing Multiple Extensible Attributes in Search Results

You can edit one or multiple extensible attributes of the matching objects in the search Results table using the *Multi-Select Edit Extensible Attributes* editor. When you change multiple extensible attribute values for selected objects, the values of all selected extensible attributes will be updated.

To edit multiple extensible attributes:

1. In the Results table, select the object checkbox. You can edit multiple extensible attribute values.
2. Click the Extensible Attributes icon.
3. In the *Multi-Select Edit Extensible Attributes* editor, click on the **Value** column to edit the value of the respective extensible attribute. For information about which values you can edit-Editing Multiple Extensible Attribute Values, see [Managing Extensible Attributes](#).

## Using the Go To Function

You can use the **Go To** function to quickly locate an object, such as a network or a DNS zone. With the autocomplete feature, you can just type the first few characters of an object name in the **Go to** field and select the object from a list of possible matches. You can also enter the entire object name, and then click **Go** to locate a specific object.

To use the Go to function:

1. From a selector, enter the first few characters of the object name in the **Go to** field. Grid Manager displays up to ten possible matches in a drop-down list.
2. Click the object from the drop-down list or use the up and down arrow keys to select the object and then press **Enter**. Grid Manager completes the operation based on the selected object.

## Opening Technical Support Requests

When you encounter product issues or require assistance, you can send a request to Infoblox Technical Support by opening a support case through Grid Manager. When you submit a support case, product information such as software version and serial number is automatically collected from the NIOS appliance on which you create the support case. You should however provide detailed information about the issue or your request, business impact, and contact information to ensure that the support request is being addressed by the appropriate resources in a timely manner. When you submit a support request, Infoblox Technical Support automatically authenticates and authorizes the contact email address that you use. It sends a confirmation email to the contact email address if the email address is registered on the Infoblox Technical Support server. If the authentication fails, you will receive an email.



 **Note**

Make sure that you enable DNS Resolver or Use SMTP Relay to create a support case from the GUI. For information about enabling DNS resolution and notifying administrators, see [Enabling DNS Resolution](#) and [Notifying Administrators](#).

Complete the following to create a support case:

 **Note**

Click **Go to the Editor** and enable **DNS Resolver** or **Use SMTP Relay** if you have not already enabled either one of them.

1. **Admin:** At the top right corner of the navigation bar, click the Admin name and select **Open a Support Case** from the drop-down menu.  
**Support:** From the Help panel -> click **Support** -> **Open a Support Case**.
2. In the *Open Support Case* editor, complete the following:
  - **To Email Address:** Enter an email address to access Infoblox Technical support. The default is [support@infoblox.com](mailto:support@infoblox.com). If you change the default email address, the email is sent to the updated email address instead of Infoblox Technical Support. Sending the email to your own email address allows you to verify and alter the email content before you forward the email to [support@infoblox.com](mailto:support@infoblox.com) or another email address.
  - **Contact Email Address:** Enter your email address or another contact email to which a confirmation is sent when the support case is created. Ensure that the contact email is legitimate and approved by Infoblox as this contact becomes the primary contact for the support case.
  - **CC Email Addresses:** Click the Add icon to add additional email addresses. This is optional.
  - **CaseType:** From the drop-down menu, select the case type:
    - **Administrative Issue:** Select this if the request is related to administrative issues.
    - **Administrative Question:** Select this if you have administrative questions.
    - **Product Issue:** Select this if the request is related to product issues.
    - **Product Question:** Select this if you have questions about Infoblox products.
  - **Severity:** Select how urgent or severe this request affects your business.
    - **Low:** Select this if the business impact is low.
    - **Medium:** Select this if the issue has moderate impact on your business.
    - **High:** Select this if this issue is urgent and requires immediate attention.
  - **Subject:** Enter a subject line for your support case.
  - **Description:** Enter a detailed description about the issue.
  - **Attach a file:** You can attach a file that contains additional information about your support case. Relevant information can help the Infoblox Support team to identify problems and troubleshoot issues in a more efficient manner.
3. Click **Send&Close** to create the support case. Optionally, click **Send&New** to send the current request and then create a new support request.

 **Note**

It might take up to 15 minutes before you get an email confirmation.

## Dashboards

The Dashboard is your home page on Grid Manager and provides easy access to tasks and a quick overview to the status of your Grid and DNS, DHCP and IPAM services. It provides easy access to tasks and a quick view to the status of your Grid and core network services. Grid Manager provides the following dashboards:

- **Tasks:** The Tasks dashboard contains task packs that provide easy access to commonly performed tasks. A task pack is a collection of tasks that belong to a specific service or function, such as IPAM or Automation. For more information, see [Tasks Dashboard](#).

- **Status:** A status dashboard contains widgets from which you can view and manage DNS, DHCP, and IPAM status and data. You can configure multiple status dashboards for managing a large number of Grid members. For more information, see [Status Dashboard](#).
- **Reporting Clustering Status:** The Reporting Clustering Status dashboard displays the reporting clustering status. This tab is displayed only if you have configured reporting clustering. For more information, see [Report Clustering Dashboard](#)

When you first log in to Grid Manager, the tasks dashboard is your home page. You can change your home page for subsequent logins.

To change your home page:

1. Navigate to any tab in Grid Manager (except for the **Dashboards** tab).
2. Click **User Profile** from the Toolbar and complete the following In the *User Profile* dialog box:  
**Default Dashboard:** Select **Status** or **Task** from the drop-down list.
3. Save the configuration.

Grid Manager displays the selected dashboard as your home page when you log in the next time.

## Tasks Dashboard

The Tasks Dashboard provides easy access to commonly performed tasks, such as adding networks and adding host records. Tasks are grouped by service-specific task packs. You must install valid licenses on the appliance to see and perform specific tasks on the Tasks Dashboard. For information about the required licenses for IPAM tasks, see [IPAM Task Pack](#).

You must also have at least read-only permission to a task-related object to add or hide the task in its task pack. To execute a task, you must have the appropriate permissions to the member and objects that are related to the tasks. For example, to add a host record from the Tasks Dashboard, you must have at least read-only permission to the host records task and read/write permission to the zone and network in which the host records are created. For information about permissions, see [Administrative Permissions for Common Tasks](#).

## Task Packs

Grid Manager displays task packs, including the IPAM and NetMRI task packs, based on valid licenses installed on the appliance. To access the IPAM task pack, you must have valid DNS or DHCP license installed on the NIOS appliance. To access the Automation task pack, you must first set up an Infoblox NetMRI appliance, install the Automation Change Management license on the NIOS appliance, and register as a user. For information about how to activate the Automation task pack, refer to the *Infoblox NetMRI Administrator Guide*.



### Note

The Tasks Dashboard will not appear in the NIOS system if no task packs are licensed for the system. Some task packs will also have dependencies. For example, the NetMRI Task Pack licensing activates along with either the MS license or the NIOS DHCP/DNS combination license. Should either of those licenses be disabled for any reason, the NetMRI Tasks will also be disabled.

To use the Automation Task Pack, you must enable the NetMRI Tasks feature set and establish a working connection between the NIOS appliance and an Infoblox NetMRI appliance. See [Enabling NetMRI Tasks](#) for details. Each task in a task pack opens a workflow dialog in which you can create task-related objects without navigating through other tabs and editors in Grid Manager. Depending on the task you perform, Grid Manager displays task results in the Result page from which you can access newly created objects, such as networks and host records. Note that when a task takes longer than usual to complete, it becomes a long running task. For information about long running tasks, see [About Long Running Tasks](#).

With valid licenses and proper registrations, Grid Manager displays the following task packs in the Tasks Dashboard:

- [Status Dashboard](#)

- [IPAM Task Pack](#)
- [NetMRI Task Pack](#)

## Status Dashboard

A status dashboard contains widgets from which you can view and manage data. Widgets are the building blocks of status dashboards. For more information about widgets, see [Adding Widgets to Dashboards](#) below. They provide information about different aspects of your Grid and networks. For example, the *Member Status* widget provides general information about a Grid member, and the *Network Statistics* widget provides data for a specified network.

The appliance provides a default status dashboard. Grid Manager displays the default dashboard only when there is more than one widget on the dashboard. You can add and modify widgets in the default dashboard, but you cannot rename or delete it. From a dashboard, you can access your most commonly accessed tasks and monitor appliance status. You can configure your own status dashboards to which you can add widgets that help you manage different data. Configuring multiple status dashboards helps organize widgets in a meaningful way and improves dashboard and widget performance. This is especially useful when you have a Grid serving a large number of Grid members. When you configure a new dashboard, you can use the existing dashboard as a template. You can create up to 100 copies at a time using the **Add Dashboard** option. For information about how to add status dashboards, see [Adding Status Dashboards](#) below.

You can add widgets to different dashboards, however, you can add only one widget at a time on each dashboard. The default number of widgets per dashboard is 10. The maximum number of widgets that you can add on each dashboard is 20 at a time. You can define the number of widgets that can be configured on each dashboard in **User Profile**. This limitation applies only to dashboards that you configure and does not apply to the default dashboard. For information about how to specify the widget limit, see [Configuring Widget Limit per Dashboard](#) below.

Grid Manager provides a default Security dashboard if you have installed any or all of the following licenses on the appliance: **Threat Protection**, **RPZ**, and **Threat Analytics**. The Security dashboard contains widgets that help you monitor the security status of the Grid. In the Security dashboard, you can add and remove widgets, but you cannot rename or delete them.

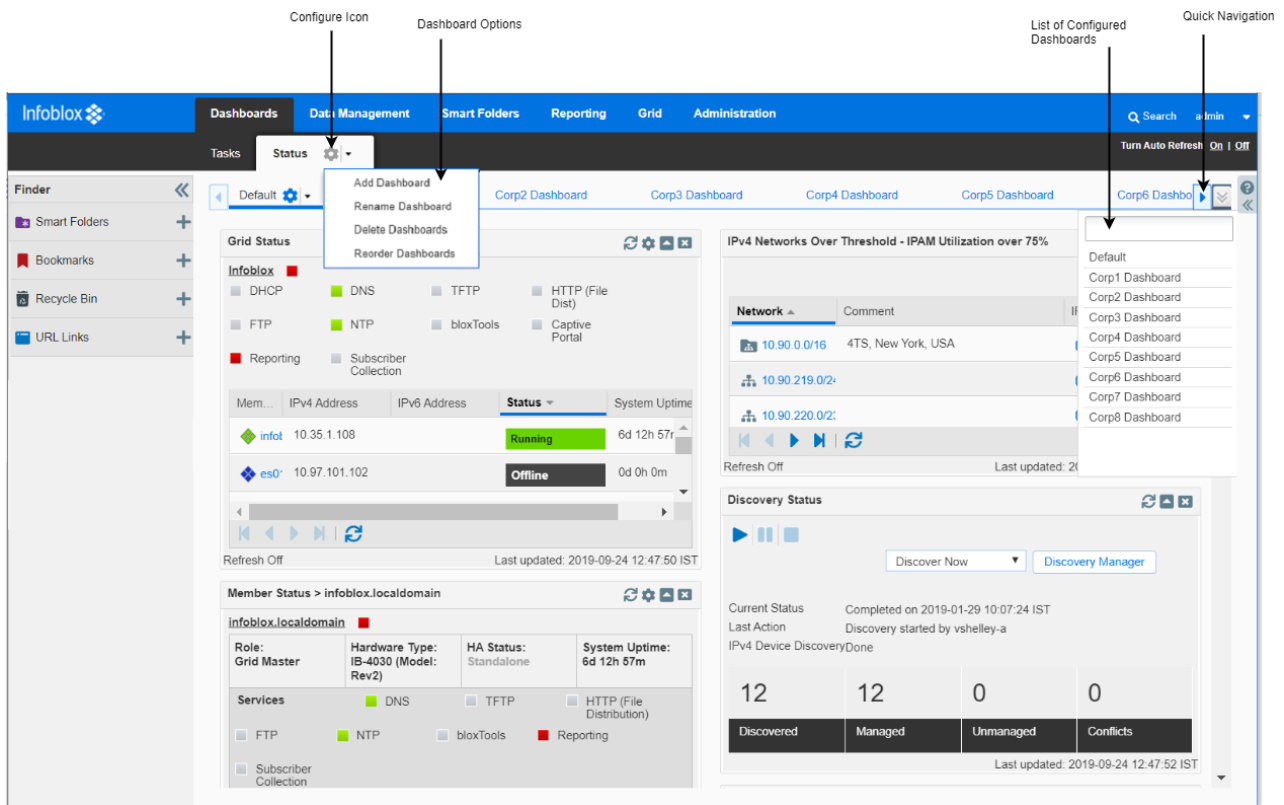


### Note

To ensure that the Security dashboard displays correct data, use NTP to synchronize the time of the Grid members with that of the Grid Master.

If you have configured a lot of status dashboards, you can use the Quick Navigation icon to quickly access each status dashboard. For information, see [Using Quick Navigation](#) below. The Status Dashboard figure below illustrates the typical layout in Grid Manager after you configure multiple status dashboards.

### *Status Dashboard*



Configure Icon Dashboard Options List of Configured Dashboards Quick Navigation

You can do the following on the **Status** tab:

- Add new status dashboards, as described in [Adding Status Dashboards](#) below.
- Rename a dashboard, as described in [Renaming Status Dashboards](#) below.
- Copy or move a widget, as described in [Copying or Moving Widgets](#) below.
- Reorder dashboards, as described in [Reordering Status Dashboards](#) below.
- Delete dashboards, as described in [Deleting Status Dashboards](#) below.
- Configure widget limit, as described in [Configuring Widget Limit per Dashboard](#) below.
- Configure Security dashboard properties, as described in [Configuring Security Status Thresholds](#) below. Adding Widgets to Dashboards

You can add all or some of the following widgets to your status dashboards depending on whether you are managing a Grid, an independent appliance, or an Infoblox Orchestration server:

- Grid Status
- Grid Upgrade Status
- Member Status (System Status)
- DNS Statistics
- Ranges Over Threshold
- IPv4 Failover Associations Status
- DHCP Statistics
- Network Statistics
- IPv4 Networks Over Threshold
- Discovery Status
- Advanced Discovery Status
- My Commands
- DDNS Statistics
- System Activity Monitor






- File Distribution Statistics
- Active WebUI Users
- Microsoft Servers Status Widget
- CSV Import Manager
- Pending Approvals
- Infoblox Community
- Mobile Devices Status
- Threat Protection Status for Grid
- Threat Protection Status for Member
- DNS Integrity Check
- *Cloud Statistics*
- Security Status for Grid
- Security Status for All Members
- Pool Licenses Statistics
- DNS Record Scavenging

Grid Manager displays the Security dashboard if you have any or all of the following licenses installed on your appliance: **Threat Protection**, **RPZ**, and **Threat Analytics**. The Security dashboard contains the following widgets, depending on the licenses installed on your appliance:

- Security Status for Grid
- Security Status for All Members
- Threat Protection Status for Grid
- Threat Protection Status for Member
- Response Policy Zone (RPZ) Status for Grid
- Response Policy Zone (RPZ) Status for Member

Note that you must have at least read-only permission to the objects that a widget displays. Otherwise, though you are allowed to select and place the widget on the dashboard, it does not display any information.

To add widgets to your dashboard:

1. **Default Status Dashboard:** From the **Dashboards** -> **Status** tab -> **Default** tab, click the Configure icon -> **Add Content**. This is applicable when you have the default dashboard only.  
**Configured Status Dashboards:** From the **Dashboards** -> **Status** tab, select the configured status dashboard, click the Configure icon -> **Add Content**.  
**Security Status Dashboard:** From the **Dashboards** -> **Status** tab -> **Security** tab, click the Configure icon -> **Add Content**. This is applicable only when at least one member in the Grid has Threat Protection, RPZ, or Threat Analytics license. Note that the Security Status dashboard is a default dashboard and it cannot be renamed or deleted.  
Grid Manager displays thumbnails of the available widgets. Use the scroll bar on the right to scroll through the widgets, as illustrated in the below Widgets Panel figure.
2. Click an icon on the filter panel, as illustrated in the Widgets Panel figure, to add a widget to the desired dashboard. The Filter panel is categorized into the following:
  - Cloud** 
  - Security** 
  - DNS/DHCP** 
  - Reset** 
When you click on an icon, Grid Manager displays thumbnails of the widgets belonging to the respective filter. If you click filters one after the other without clicking **Reset**, Grid Manager displays thumbnails of all widgets along with the icon that indicates the category to which the widget belongs. Click **Reset** to view only those widgets that belong to the selected category.
3. Select and drag a widget to the desired location on your dashboard. You can also click  icon to add a widget to the desired dashboard.  
After you add a widget to the dashboard, you can configure it to provide relevant data. You can also copy or move a widget, by selecting and dragging it to its new location on your dashboard. Grid Manager saves your dashboard configuration and displays it the next time you log in.  
You can turn on auto-refresh by clicking **On** in the **Turn Auto Refresh** field at the top of the dashboard to

periodically refresh the contents of all widgets in the dashboard. Click **Off** to disable auto-refresh for all widgets in the dashboard. When auto-refresh is disabled, you can enable it for individual widgets by clicking the Configure icon in the corresponding widgets. You can specify the auto-refresh period in seconds. The default auto-refresh period is 30 seconds.

**Warning**

If the *Detailed Status panel* is open, the following actions take place:

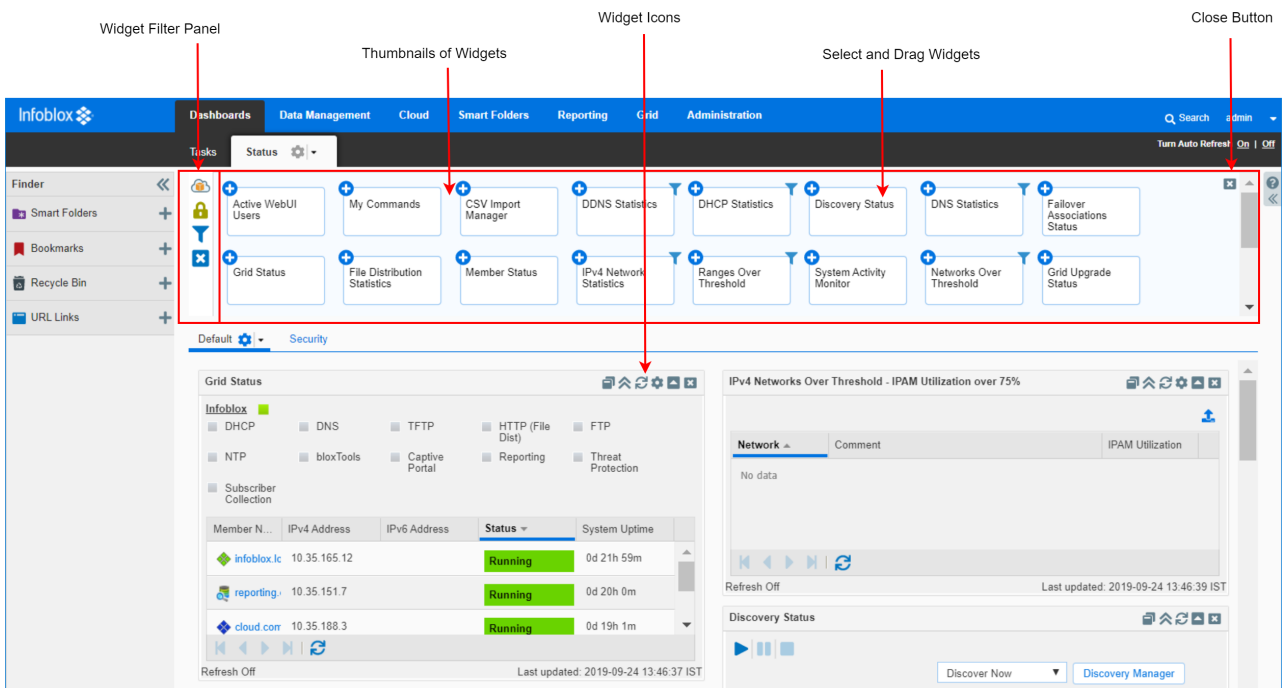
- Grid Manager auto refreshes at a rate of 30 seconds.
- Widgets that support user-specified auto refresh, refresh at the rate specified in the **Auto Refresh Period** field.

Therefore, even if you set the session timeout to be to greater than the auto refresh rate, auto refresh still takes place. The Grid Manager session does not time out because auto refresh takes precedence over the session timeout.

Widgets have the following icons:

- **Copy/Move:** Click to copy or move the widget from a dashboard to another. For information about how to copy or move, see Copying or Moving Widgets below.
- **Span Up/Span Down:** Click to resize the widget. Click **Span Up** to increase the width of the widget. Click **Span Down** to decrease the width of the widget. Note that the fully spanned widgets are moved to the top of the dashboard.
- **Refresh:** Click to update the content of the widget. Each widget contains a status bar at the bottom that displays the last date and time it was updated.
- **Configure:** Click to hide and show the configuration options of the widget.
- **Toggle:** Click to minimize and restore the widget.
- **Close:** Click to remove the widget from a dashboard.

### Widgets Panel



### Configuring Widget Limit per Dashboard

You can define the number of widgets that can be configured on each dashboard. This limitation applies only to dashboards that you configure and does not apply to the default dashboard.

1. From the **Dashboards** -> **Status** tab, click the Configure icon -> **User Profile**.
2. In the *User Profile* editor, complete the following:  
**Maximum Widgets per Dashboard:** Specify the maximum number of widgets that can be configured per Dashboard. You can enter a value between 1 and 20. The default value is 10. This limit does not apply to the default dashboard.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Adding Status Dashboards

You can create your own status dashboards and add the widgets that you need. You can configure up to 100 status dashboards at a time. When you create multiple instances of a dashboard, the appliance names each dashboard by adding an incremental suffix to the name of the new dashboard. For example, if you name a new dashboard "Corp\_Dashboard" and specify the number of instances as three, then the appliance creates three instances of this new dashboard. In this example, the appliance creates three dashboards: Corp\_Dashboard, Corp\_Dashboard1, and Corp\_Dashboard2. Note that the dashboards you create will not be available to other users. You cannot share dashboards you have created with other users.

Note that "Security" is reserved for the default Security dashboard. Grid Manager displays an error message if you name a new dashboard "Security".

To add a new status dashboard, complete the following:

1. From the **Dashboards** -> **Status** tab, click the Configure icon -> **Add Dashboard**.
2. In the *Add Dashboard* wizard, complete the following:
  - **Name:** Enter a name for the new dashboard.
  - **Add <> instances of this new dashboard:** Enter the number of dashboards you want to create. The maximum number of dashboards you can create is 100 at a time.
  - **Copy initial content from an existing dashboard:** Select this checkbox if you want the appliance to copy the contents from an existing status dashboard into the new dashboard. After you select this checkbox, the appliance displays the list of configured dashboards. Select a dashboard from the list. By default, this checkbox is not selected.
3. Save the configuration.

The appliance displays all dashboard instances on the **Status** tab.

## Using Quick Navigation

You can use the Quick Navigation icon to quickly access a specific dashboard. The appliance provides the Quick Navigation icon at the right corner of the status dashboards, as illustrated in the *Status Dashboard* figure above.

To quickly navigate to a dashboard, complete the following:

1. From the **Dashboards** -> **Status** tab, click the Quick Navigation icon at the right corner of the dashboards. The list of configured dashboards is displayed.
2. Select a dashboard or specify the name of the dashboard in the text box. The appliance displays the selected dashboard.

## Renaming Status Dashboards

You can rename only the status dashboards that you have configured. You cannot rename the default dashboard and the Security dashboard.

To rename a dashboard, complete the following:

1. From the **Dashboards** -> **Status** tab, click the Configure icon -> **Rename Dashboard**.
2. In the *Rename Dashboard* wizard, complete the following:
  - **Select a dashboard:** Select a dashboard from the drop-down list.
  - **Name:** Enter the new name of the dashboard.
3. Do one of the following:

- Click **Save and Close** to save the new name and close the wizard.
- Click **Save** to save the new name and continue to rename other dashboards.

To rename a specific dashboard:

1. From the **Dashboards** -> **Status** tab, select a dashboard that you want to rename.
2. Click the Configure icon -> **Rename Dashboard**.
3. In the *Rename Dashboard* wizard, enter the new name in the **Name** text box.
4. Click **Save and Close** to save the new name and close the wizard.

## Copying or Moving Widgets

You can copy or move a widget from one dashboard to another. When you add a widget that already exists, the appliance displays an error message. When you move a widget, it is moved from the source to the destination dashboard. The moved widget will not be available in the source dashboard anymore. When you copy a widget, the widget is duplicated and is available in both the source and destination dashboards. Note that the Copy/Move icon is not available in a widget if the appliance has only the default status dashboard.

To move or copy a widget, complete the following:

1. From the **Dashboards** -> **Status** tab, select a status dashboard.
2. Select the widget that you want to copy or move, and then click the Copy/Move icon.
3. In the *Copy/Move <name of the widget>* wizard, complete the following:
  - **Copy**: Select this to copy a widget.
  - **Move**: Select this to move a widget.
  - **To Dashboard**: Select the name of the destination dashboard.
4. Click **OK**.

## Reordering Status Dashboards

You can change the order of your status dashboards. When you add a new status dashboard, it is added as a tab. When you create multiple instances of a dashboard, they are added as subsequent tabs. You can arrange the order of each dashboard through the reordering process.

To reorder status dashboards, complete the following:

1. From the **Dashboards** -> **Status** tab, click the Configure icon -> **Reorder Dashboards**.
2. The following are displayed in the *Order Dashboards* wizard:
  - **Ordering**: You can use the up and down arrows to move dashboards in the desired order or drag and drop them to the desired positions.
  - **Dashboard**: Displays the list of all the status dashboards.
3. Click **OK** to save the changes.

## Deleting Status Dashboards

You can delete status dashboards that you have configured. You cannot delete the default status dashboard and the Security dashboard. You can delete multiple dashboards at the same time. Note that you cannot restore a deleted dashboard.

To delete multiple dashboards, complete the following:

1. From the **Dashboards** -> **Status** tab, click the Configure icon -> **Delete Dashboards**.
2. In the *Delete Dashboards* wizard, select the **Dashboard** checkbox. You can select multiple checkboxes for multiple dashboards.
3. Click **Delete**.
4. Click **Yes** in the *confirmation* dialog box. To delete a specific dashboard:
5. From the **Dashboards** -> **Status** tab -> select the *<Status Dashboard>* tab.
6. Click the Configure icon -> **Delete Dashboard**.
7. *In the Delete Confirmation* dialog box, click **Yes**.



## Configuring Security Status Thresholds

You can configure thresholds to determine the overall status of Threat Protection, DNS RPZ (Response Policy Zone), and DNS Threat Analytics services in the Grid. Grid Manager provides a view of the overall security status of the Grid in the *Security Status for Grid* dashboard widget. For information, see *Security Status for Grid* below.

To configure the thresholds for security status, complete the following:

1. From the **Dashboards** -> **Status** tab, click the Configure icon -> **Global Dashboard Properties**.
2. In the *Global Dashboard Properties* editor, complete the following:
  - **Threat Protection Thresholds:** Define the thresholds for each severity level of the threat protection events for the following colors:
    - **Yellow:** Specify the low threshold value for **Critical**, **Major**, and **Warning** severity level. The default values are 1, 20, and 100 for **Critical**, **Major**, and **Warning** respectively.
    - **Red:** Specify the high threshold value for **Critical**, **Major**, and **Warning** severity level. The default values are 5, 100, and 1000 for **Critical**, **Major**, and **Warning** respectively.

Depending on the specified thresholds, Grid Manager determines the status of threat protection service as follows, which is displayed in the **Status** column of the *Security Status for Grid* widget:

- **Green (OK):** When the number of threat protection events are less than the low threshold value specified for the yellow color for all the severity levels.
  - **Yellow (Warning):** When the number of threat protection events equals or exceeds the threshold value specified for the yellow color but less than the threshold value specified for the red color for any of the severity levels.
  - **Red (Critical):** When the number of threat protection events equals or exceeds the high threshold value specified for the red color for any of the severity levels.
- **Response Policy Zone Thresholds:** Define the threshold values for the following colors to determine the overall status of RPZ:
    - **Yellow:** Specify the low threshold value for **Blocked**, **Substitute**, and **Passthru** RPZ rules. The default values are 10, 1, and 100 for **Blocked**, **Substitute**, and **Passthru** respectively.
    - **Red:** Specify the high threshold value for **Blocked**, **Substitute**, and **Passthru** RPZ rules. The default values are 100, 10, and 1000 for **Blocked**, **Substitute**, and **Passthru** respectively.

Depending on the specified thresholds, Grid Manager determines the status of RPZ as follows, which is displayed in the **Status** column of the *Grid Security Status* widget:

- **Green (OK):** When the number of RPZ hits are less than the low threshold value specified for the yellow color for all the rule types.
  - **Yellow (Warning):** When the number of RPZ hits equals or exceeds the threshold value specified for the yellow color but less than the threshold value specified for the red color for any of the rule types.
  - **Red (Critical):** When the number of RPZ hits equals or exceeds the high threshold value specified for the red color for any of the rule types.
- **Threat Analytics Thresholds:** Define the thresholds for the following colors, to determine the overall status of DNS Threat Analytics:
    - **Yellow:** Specify the low threshold value for DNS Tunneling events. The default value is 1.
    - **Red:** Specify the high threshold value for DNS Tunneling events. The default value is 5.

Depending on the specified thresholds, Grid Manager determines the status of DNS Threat Analytics as follows, which is displayed in the **Status** column of the *Grid Security Status* widget:

- **Green(OK):** When the number of DNS tunneling attacks are less than the low threshold value specified for the yellow color.
  - **Yellow(Warning):** When the number of DNS tunneling attacks equals or exceeds the threshold value specified for the yellow color but less than the threshold value specified for the red color.
  - **Red(Critical):** When the number of DNS tunneling attacks equals or exceeds the high threshold value specified for the red color.
3. Save the configuration.

 **Note**

If you have configured the same threshold value for both Yellow and Red color in the *Global Dashboard Properties* editor and if the same number of events are triggered, then Grid Manager displays the status in red in the *Grid Security Status* widget.

## Grid Status

The *Grid Status* widget provides status information about the Grid members and services. Add the *Grid Status* widget to your Dashboard to monitor the Grid status.

You can configure the *Grid Status* widget to display information about all Grid members or only Grid members that have service errors. To modify the *Grid Status* widget, click the Configure icon and select one of the following:

- **Show all Grid members** (this is the default)
- **Only show members with service warnings or errors:** When you select **Only show members with service warnings or errors**, the widget displays only the members that have service errors. The widget does not display any data in the member table if all the services on all members are running properly.
- **Group Members by:** If you want to group members by the same extensible attribute value, select this and choose an extensible attribute from the drop-down list. The appliance groups Grid members that have the same extensible attribute value, and the Grid Status displays the following information:
  - **<Extensible Attribute Name>:** The value of the selected extensible attribute. You can click the link of the extensible attribute value to view all the members in this group in the Grid/Members view.
  - **Status:** This is the overall status for all members in the group. Depending on the status of each member, the overall status can be one of the following:
    - **Working:** Indicates that all the members in the group are running properly.
    - **Warning:** Indicates that one of the members in the group has operational problems. For example, if there are two members in a group with one member **Running** and another member is **Offline**, then the overall status will be **Warning**.
    - **Failed:** Indicates that at least one of the members in the group is in the failed status and none of the members in the group are in the **Running** or **Working** status. For example, if there are two members in the group and one of them is in **Failed** status and the other is **Offline**, then the overall status is **Failed**.
    - **Offline:** Indicates that one or more members in the group is offline and none of the members in the group are in the **Failed** or **Running** status. For example, if a member is in the **Working** status and another member is in the **Offline** status, then overall status is **Offline**.
    - **Inactive:** Indicates that one or more members in the group is inactive and none of the members in the group are in the **Failed**, **Offline**, **Working**, or **Running** status.
    - **Unknown:** Indicates that the status of all the members in the group is unknown.




 **Note**

You can click a member link to monitor the detailed status of the selected member. Grid Manager displays the **Grid** tab -> **Member** tab. You can click on a group to show the members of the group in the Grid/Members view.





The *Grid Status* widget also displays the following information in the member table:

- **Member Name:** The name of the member.
- **IPv4 Address:** The IPv4 address of the member.
- **IPv6 Address:** The IPv6 address of the member.
- **Status:** The current status of the member.
- **System Uptime:** The duration of time (days, hours, and minutes) that the Grid member has been up and running.

In the upper section of the widget, Grid Manager displays the overall status of the Grid. The Grid status represents the status of the most critical member in the Grid. When all Grid members are running properly, the overall Grid status is green. When one of the members has operational issues, the overall Grid status is red. The status icon can be one of the following:

Icon	Color	Meaning
	Green	All Grid members are operating normally in a "Running" state.
	Yellow	At least one of the Grid members is connecting or synchronizing with its Grid Master.
	Red	At least one of the Grid members does not have a Grid license, is offline, upgrading, downgrading, or shutting down.

This section also displays the overall operational status of the DNS, DHCP, NTP, FTP, TFTP, HTTP (File Distribution), bloxTools, Captive Portal, Subscriber Collection, Cloud DNS Sync, DNS Accelerator usage, and Reporting services that are currently running on the Grid. The status icon can be one of the following:

Icon and Color	Meaning
 Green	The enabled service is running properly on one or more Grid members.
 Yellow	At least one of the Grid members is having issues with the enabled service.
 Red	The enabled service is not running properly on at least one of the members. (A red status icon can also appear temporarily when the service is enabled and begins running, but the monitoring mechanism has not yet notified Grid Manager.)
 Gray	<b>The service is not configured or is disabled on at least one Grid member.</b>

## Grid Upgrade Status

The *Grid Upgrade Status* widget provides upgrade status of the Grid Master and members. Add the *Grid Upgrade Status* widget to your Dashboard to monitor the upgrade status of the Grid and its members.

The *Grid Upgrade Status* widget displays the following information:

- **Upgrade Status:** The current upgrade status of the Grid. This can be **Running**, **Paused**, **Canceled**, or **Inactive**.
- **Grid Member Upgrade Process Status:** The pie chart shows the number of members that are still processing the upgrade, members that have completed the upgrade, and members that are waiting for the upgrade to happen.
- **Detailed Upgrade Status:** Click this link to access the **Grid** tab -> **Upgrade** tab to see detailed information about the upgrade.

The table on the right shows a summary of the upgrade status of the upgrade groups. It displays the following information:

- **Group:** The name of the upgrade group.
- **Date/Time:** The date and time when the upgrade started on this upgrade group. Note that the time zone is the time zone of the first member in the upgrade group.
- **Completed:** Indicates whether the upgrade is complete or not.

## Member Status (System Status)

The *Member Status* widget provides status information about the system resources and services of a Grid member, including the reporting server. The *System Status* widget provides the operational status about an independent appliance. Add a *Member Status* widget to your Dashboard for each Grid member that you want to monitor. The widget

always displays the services that a Grid member is running. You can then configure it to display additional information and specify how the information is displayed.




You can modify the *Member Status* or the *System Status* widget by clicking the Configure icon. If you have an independent appliance, you can only configure some of the following:


- For *Member Status* widget only: Click **Select Member** to select a Grid member for display. When you select the reporting server, the widget displays reporting usage.
- Select the information you want to display:
  - **Show Role:** For *Member Status* widget only. Click to display whether the appliance is a Grid Master, Grid Master candidate, or Grid member. An independent appliance does not have a Grid license installed.
  - **Show Hardware Type:** Click to display the appliance hardware model.
  - **Show HA Status:** Click to display whether the appliance is part of an HA pair. It displays one of the following:
    - **Standalone:** The Grid member is an independent appliance.
    - **HA OK:** The Grid member is part of an HA pair that is functioning properly.
    - **HA Broken:** The appliance is part of an HA pair that is not operating properly. You can check the logs to determine the problem.
  - **Show System Uptime:** Click to display the duration of time (days, hours, and minutes) that the Grid member has been up and running.
- **Statistics:** Select the data that you want to display and its format:
  - **CPU:** Click to display the percentage of CPU that is in use. Select either **Dial** or **Bar** for the display format.
  - **Memory:** Click to display the current percentage of memory that is in use. Select either **Dial** or **Bar** for the display format.
  - **Database:** Click to display the percentage of the database that is in use. Select either **Pie** or **Bar** for the display format.
  - **Disk:** Click to display the percentage of the data partition on the hard disk drive in use. Select either **Pie** or **Bar** for the display format.
  - **System Temperature:** Click to display the system temperature. Depending on the hardware model, the system temperature may not be available. Select to display the temperature in either **Celsius** or **Fahrenheit**.
  - **CPU Temperature:** Click to display the CPU temperature. Depending on the hardware model, the CPU temperature may not be available. Select to display the temperature in either **Celsius** or **Fahrenheit**.

Click the Configuration icon again to hide the configuration panel after you complete the modification.

Grid Manager displays the hostname of the appliance at the top of the widget. You can click the name link to view detailed information about the appliance. The widget also displays the upgrade status if the member is currently in the process of an upgrade. If the member is scheduled for an upgrade, the **Scheduled for upgrade** link appears. You can click this link to access the **Grid** tab -> **Upgrade** tab to view more details about the date and time of the scheduled upgrade.

The widget also displays the service status of the following: FTP, TFTP, HTTP (File Distribution), DNS, DHCP, NTP, bloxTools, Captive Portal, Subscriber Collection, Cloud DNS Sync, DNS Accelerator, and Reporting in the *Services* section. The service status can be one of the following:

Icon	Color	Meaning
	Green	The service is enabled and running properly.
	Yellow	The service is enabled, but there may be some issues that require attention.
	Red	The service is enabled, but it is not running properly or is out of synchronization. (A red status icon can also appear temporarily when a service is enabled and begins running, but the monitoring mechanism has not yet notified the GUI engine.)

Icon	Color	Meaning
	Gray	The service is not configured or is disabled.

The widget also displays the statistics you specified, such as CPU usage, memory, and database usage, in the format you selected.

When you select the reporting server, you can also see the reporting usage information:

- **Reporting Usage:** Displays the daily consumption rate for the reporting service.

For more information about reporting, see [Infoblox Reporting and Analytics](#).

## DNS Statistics

The *DNS Statistics* widget provides statistics for a member or for a zone. The zone statistics are cumulative, collected from all the members that are authoritative servers for zones or are hosting stub zones. The widget displays the totals for each type of DNS response as well as a line graph that tracks the responses per second.

You can add a *DNS Statistics* widget to your Dashboard for each zone or member DNS server on the Grid. To configure the *DNS Statistics* widget, click the Configure icon and do the following:

- Click **Select Member**. In the *Member Selector* dialog box, choose a Grid member to display statistics for all its stub zones and authoritative zones.  
or
- Click **Select Zone**. In the *Zone Selector* dialog box, choose a DNS zone to display statistics for that zone only.

The widget displays only the option that you selected on your subsequent logins. For example, if you clicked **Select Member**, the widget displays the **Select Member** option only, and not the **Select Zone** option, when you log in again.

- **Graph Configuration:** Select which DNS messages you want to track in the **Responses per Second** graph.
  - **Success:** The number of successful queries.
  - **NXDOMAIN:** The number of queries for domain names that did not exist in the database.
  - **Referral:** The number of queries that became referrals.
  - **NXRRSET:** The number of queries for domain names that did not have the requested records.
  - **Failure:** The number of queries that failed due to reasons other than nonexistent domain names or records in a domain.
  - **Recursion:** The number of recursive queries for which the name server sent queries to other name servers.

The widget displays the following information:

- **DNS Responses** tab: Displays a pie chart and the total number of each type of message. It also displays the total number of full and incremental zone transfers that the Grid member performed.
- **Responses per Second** tab: Displays a line graph that tracks the DNS responses received per second, within an hour. The time is displayed according to the time zone specified in the User Profile. If the auto-detect time zone option is enabled and Grid Manager cannot determine the browser time zone, then the time is displayed in UTC format. You can mouse over the graph to display the coordinates of any point in the graph.

## Ranges Over Threshold

The *Ranges Over Threshold* widget enables you to monitor IPv4 DHCP range usage from your Dashboard. It lists the IPv4 ranges that are allocated above a specified threshold and thus may warrant your attention. The default threshold is 75%. For information, see [Configuring Thresholds for DHCP Ranges](#). Note that the appliance highlights disabled IPv4 ranges in gray.

The widget displays the IPv4 ranges with utilization percentages that surpass the threshold. To configure the *Ranges Over Threshold* widget, click the Configure icon and do the following:

- **Network View:** Select a network view in which you want to monitor the IPv4 ranges. This field is displayed only when you have more than one network view.
- **Threshold:** Enter a new threshold value. The default is 75%.  
In addition, you can do the following:
  - Click the Export button to export the list of IPv4 ranges that surpass the threshold to a file in CSV format.
  - Click the Refresh button to refresh the data in the list.

## IPv4 Failover Associations Status

The *IPv4 Failover Associations Status* widget enables you to monitor the status of the failover associations from your Dashboard. It lists all the failover associations in the Grid and displays their names and status. The widget also displays the primary and secondary servers in the association. When you click a failover association link or a status link, Grid Manager displays the Failover Association section where you can get detailed information about the failover association. For information, see [Managing Failover Associations](#).

In addition, you can do the following:

- Click the Export button to export the list of failover associations to a file in CSV format.
- Click the Refresh button to refresh the data in the list.

## DHCP Statistics

The *DHCP Statistics* widget displays statistics about the different types of DHCP messages that a Grid member sends and receives. The widget displays the totals for each type of DHCP message as well as a line graph that tracks the messages per second.

You can add a *DHCP Statistics* widget to your Dashboard for each member DHCP server in the Grid. If the DHCP service is not enabled or is offline, the widget displays a message indicating that the DHCP statistics are not available.

To configure the *DHCP Statistics* widget, click the Configure icon and do the following:

- **Protocol:** Select either **IPv4** or **IPv6**.
- Click **Select Member**. In the *Member Selector* dialog box, select a Grid member from the list.
- **Graph Configuration:** This section lists IPv4 or IPv6 messages, depending on the protocol you selected.
- Select which IPv4 messages you want to track in the **Messages per Second** graph.
  - **Discovers:** The number of DHCPDISCOVER messages that the Grid member received from DHCP clients. A DHCP client broadcasts a DHCPDISCOVER message to obtain an IP address.
  - **Offers:** The number of DHCPOFFER messages that the Grid member sent to DHCP clients. If the Grid member has an IP address that it can allocate to the DHCP client that sent the DHCPDISCOVER message, the Grid member responds with a DHCPOFFER message that includes the IP address and configuration information.
  - **Requests:** The number of DHCPREQUEST messages that the Grid member received from DHCP clients. A DHCP client sends DHCPREQUEST messages when it selects a lease, connects to the network, and if it renews the lease.
  - **Acks:** The number of DHCPACK messages that the Grid member sent to DHCP clients. When the Grid member receives a DHCPREQUEST message, it responds with a DHCPACK message to confirm the IP address selected by the DHCP client.
  - **Nacks:** The number of DHCPNACK messages that the Grid member sent to DHCP clients. The Grid member sends a DHCPNACK message when a DHCP client requests an IP address that is not valid for the network.
  - **Declines:** The number of DHCPDECLINE messages that the Grid member received. A DHCP client sends a DHCPDECLINE message to a DHCP server when it discovers that the IP address offered by a DHCP server is already in use.
  - **Inform:** The number of DHCPINFORM messages that the Grid member received. A client that did not receive its IP address from the DHCP server can send it a DHCPINFORM message to retrieve configuration parameters, such as the IP addresses of DNS servers in the network.
  - **Releases:** The number of DHCPRELEASE messages that the Grid member received. A DHCP client sends a DHCPRELEASE message when it terminates its lease and releases its IP address.
- Select which IPv6 messages you want to track in the **Messages per Second** graph.

- **Declines:** The number of Decline messages that the Grid member received. A DHCP client sends a Decline message to a DHCP server when it discovers that the IP address offered by a DHCP server is already in use.
- **Renews:** The number of Renew messages that the Grid member received. A DHCP client sends a Renew message to a DHCP server to extend the lifetimes on the leases granted by the DHCP server and to update other properties.
- **Information Requests:** The number of Information-Request messages that the Grid member received. A client sends an Information-Request message to retrieve configuration parameters, such as the IP addresses of DNS servers in the network.
- **Solicits:** The number of Solicit messages that the Grid member received, including Solicit messages embedded in Relay-Forward messages. A DHCP client sends a Solicit message to locate DHCP servers.
- **Requests:** The number of Request messages that the Grid member received. A DHCP client sends a Request message to request one or more IP addresses and configuration parameters from a DHCP server.
- **Rebinds:** The number of Rebind messages that the Grid member received. A DHCP client sends a Rebind message to extend the lifetime of its lease and to update configuration parameters.
- **Releases:** The number of Release messages that the Grid member received. A DHCP client sends a Release message when it terminates its lease and releases its IP address.
- **Advertises:** The number of Advertise messages that the Grid member sent. When a DHCP server receives a Solicit message, it can respond with an Advertise message to indicate that the server is available for DHCP service.
- **Replies:** The number of Reply messages that the Grid member sent. A DHCP server sends a Reply message that includes IP addresses and configuration parameters when it responds to Solicit, Request, Renew or Rebind message. It sends a Reply message with configuration parameters only when it responds to an Information-Request message.

The widget displays the following information:

- **DHCP Messages** tab: Displays a pie chart and the totals for each type of DHCP message. It also displays the number of Deferred Updates, which are DDNS update requests which are deferred because the DNS primary was not reachable when the update was first attempted.
- **Messages per Second** tab: Displays a line graph that tracks the DHCP messages that were sent and received per second, within an hour. The time is displayed according to the time zone specified in the User Profile. If the auto-detect time zone option is enabled and Grid Manager cannot determine the browser time zone, then the time is displayed in UTC format. You can mouse over the graph to display the coordinates of any point in the graph.

## Network Statistics

The *Network Statistics* widget provides information about IP address usage in an IPv4 network. You can monitor several networks simultaneously to view the distribution of address resources. Such information can indicate if there is a sufficient number of available addresses in each network. It can also provide information about the distribution of address resources, indicating if there are too many unused addresses in one network while all the addresses in another are in use.

Add a *Network Statistics* widget to your Dashboard for each network that you want to monitor. You can monitor IPv4 networks only.

To configure the *Network Statistics* widget, click the Configure icon and do the following:

- Select one of the following chart types:
  - **Pie**
  - **Bar**
- Click **Select Network**. In the *Network Selector* dialog box, choose a network from the list and click **Select**. Note that if multiple network views were previously configured, Grid Manager displays the default network view. You can choose another network view from the drop-down list, and then select a network. The *Network Statistics* widget displays the following information about the selected network:
  - **IPAM Utilization:** When you define a network, this is the percentage based on the IP addresses in use divided by the total addresses in the network. For example, in a /24 network, if there are 25 static IP addresses defined and a DHCP range that includes 100 addresses, the total number of IP addresses in use is 125. Of the possible 256 addresses in the network, the IPAM utilization is about 50% for this network. When you define a network container that contains subnets, this is the percentage of the total address space defined within the container



regardless of whether any of the IP addresses in the subnets are in use. For example, when you define a /16 network and then 64 /24 networks underneath it, the /16 network container is considered 25% utilized even when none of the IP addresses in the /24 networks is in use.

You can use this information to verify if there is a sufficient number of available addresses in a network. The IPAM utilization is calculated approximately every 15 minutes.

- **Unmanaged:** The number of discovered IP addresses that do not have corresponding records on the appliance, such as A records, PTR records, fixed address records, host records, or leases. To obtain this data, you must run a discovery process on the network first.
- **Conflicts:** The number of IP addresses that have either a MAC address conflict or a DHCP range conflict. To obtain this data, you must run a discovery process on the network first. A discovered host has a MAC address conflict when its MAC address is different from that specified in its fixed address, DHCP lease, or host record. A discovered host has a DHCP range conflict when it is part of a DHCP range, but it does not have a matching fixed address or DHCP lease, and it is not part of an exclusion range.

## Network Users - Active Users

The *Network Users Active Users* widget displays up to 10 active users for Windows devices managed by the Grid. To modify the *Network Users Active Users* widget, click the Configure icon and select one of the following:

- **Only show networks with:** Select **At least** or **Fewer than** from the drop-down list and specify the number of users in the **Users** field. The default values are set to **At least** and **10** users.
- **All Network Views:** Select this to monitor active users on the managed Active Directory domains in all network views.
- **Select Network View:** Click Select to select the network view in which you want to monitor active users. The *Network Users-Active Users* widget displays the following information:
  - **Network:** The network address. You can click the network link to view network details.
  - **Active Users:** All users who are currently using the Active Directory domains. You can also export the list to a .csv file.

## IPv4 Networks Over Threshold

The *IPv4 Networks Over Threshold* widget enables you to monitor IPv4 network and IP address usage from your Dashboard. It lists the IPv4 networks that are allocated above a specified threshold and thus might warrant your attention. The default threshold is 75%.

For network containers, the threshold is the percentage of IP address space that has been allocated. For subnets, it is the percentage of used addresses, except the broadcast and network addresses. The widget displays the network containers and subnets with utilization percentages that surpass the threshold.

You can also select to view IPv4 cloud networks only if you have deployed Cloud Network Automation. For information about this feature, see [Deploying Cloud Network Automation](#).

To configure the *Networks Over Threshold* widget, click the Configure icon, and then complete the following:

- **Threshold:** Enter a new threshold value. The default is 75%.
- **Type:** Select **IPAM Utilization** or **IPv4 DHCP Utilization**. For information, see [Managing IPv4 DHCP Data](#).
- **All Network Views:** Select this to monitor threshold for IPv4 networks in all network views.
- **Select Network View:** Click Select to select the network view in which you want to monitor the threshold.

To view information related to cloud networks, select **View Cloud Networks Only**, and then select one of the following:

- **All Tenants:** Displays information for all tenants.
- **Select Tenant:** Click **Select** to select a specific tenant. In addition, you can do the following in this widget:
  - Click the Export button to export the list of networks that surpass the threshold to a file in CSV format.
  - Click the Refresh button to refresh the data in the list.



## Port Status

The Port Status widget provides a quick way to inspect the interface status for any discovered device in the network. The widget shows an overview of all interfaces on all devices or for a single device (called the Data Scope).

Click the Configure icon to change settings for the widget.

1. You can choose a **Bar** or **Pie** chart for the **Total Switch or Switch-Router** chart, which shows the percentage of ports that are operationally Active and that are operationally Down.
2. Under **Data Scope**, the **All Devices** setting allows the widget to show the total counts for all discovered network infrastructure devices. This is the default.
3. To use the widget to display port information for a single device, such as a switch, enable the **Select Device** radio button. Choose the device in the *Device Selector* window. The widget adjusts its reported values to the scale of the selected device.
4. You can also choose the **Media Type**. to further filter port status information. Choices include: **Ethernet Interface**, **Layer 2 Virtual LAN**, **Proprietary Serial Interface**, **Proprietary Virtual/Internal Interface**, **Loopback Interface** and **Tunnel Interface**.

The counters in the widget include **Total Switch and Switch-Router Ports**, **Total Down Switch and Switch-Router Ports** and **Total Active Switch and Switch-Router Ports**.

## Discovery Status

The appliance can run an IP discovery to detect and obtain information about active hosts in specified networks. For information about the discovery process, see [About Discovery](#).

You can add the *Discovery Status* widget to your Dashboard. From this widget, you can access Discovery Manager and configure parameters for a discovery. You can do the following from the widget:

- Start a discovery immediately. For more information about immediate discovery, see [Configuring IP Discovery](#).
- Schedule a discovery for a later date and time. For more information about discovery, see [Configuring IP Discovery](#).
- Configure a recurring discovery. For more information about recurring discovery, see [Configuring IP Discovery](#).
- Click the Start button to start a discovery process.
- Click the Pause button to temporarily pause the process.
- Click the Stop button to stop the process.

This widget displays the status of discovery tasks. If there are no active discovery tasks, the widget displays the discovery results of the previous tasks. For information about starting and scheduling a discovery task, see [Guidelines Before Starting a Discovery](#).

After you start a discovery, the *Discovery Status* widget displays a status bar that indicates the discovery is in progress. It also tracks the number of networks in an IP discovery. You can click the Refresh icon to update the discovery status.

The widget displays the following information about the discovery process:

- **Current Status:** If a discovery is in progress, this field displays its current status. Otherwise, it displays the date and time of the last discovery.
- **Last Action:** Displays the last operation and the admin who initiated it.
- **IPv4 Device Discovery:** Displays the total number of IPv4 networks and the IPv4 network and IP address range on which the IP discovery is currently running. You can click **Refresh** to update this information.

The *Discovery Status* widget also displays the following information about the last discovery:

- **Discovered:** The total number of active hosts in the network.
- **Managed:** The number of discovered IP addresses that are managed by the NIOS appliance. These IP addresses have an A record, PTR record, fixed address record, host record, lease, or are within a configured DHCP range.
- **Unmanaged:** The number of discovered IP addresses that do not have corresponding records on the appliance, such as A records, PTR records, fixed address records, host records, or leases.
- **Conflicts:** The number of discovered hosts that have a MAC address conflict or are part of a configured DHCP range, but do not have a fixed address or lease record and are not part of an exclusion range.

## Advanced Discovery Status

With the correct licensing, dedicated NIOS appliances operating as Grid members can perform infrastructure device discovery. NIOS appliances with the Discovery license operate primarily for discovery tasks and do not perform core DNS or DHCP network functions. Discovery appliances, called Probes, collect all network device data and compile it into a database. A separate NIOS appliance, called a Consolidator, aggregates the collected device information from the Probes and synchronizes with the Infoblox Grid Master.

For more information about discovery and its features and requirements, see [Infoblox Network Insight](#) and its associated sections.

The Advanced Discovery Status widget provides several basic counts describing the general state of device discovery within the Grid, and for networks outside the Grid being inventoried by the NIOS appliances designated for discovery. The widget divides counters into two categories: **Networks** and **Assets**. Network counters refer to counts of managed and unmanaged networks discovered by Probe appliances. Asset counters refer to counts of specific types of network devices, termed Assets, which are comprised of end hosts, enterprise servers, enterprise printers, and any other enterprise asset that exists in an end-user network segment. The widget counters include:

In the **Networks** category:

- **Discovered:** The total number of networks discovered by Probe appliances.
- **Managed:** The number of discovered networks that are currently managed by the NIOS Grid. These IP networks have been converted from Unmanaged status to Managed status.
- **Unmanaged:** The number of discovered networks that are counted as Unmanaged by the NIOS Grid Master. After a network is discovered and catalogued by a Probe appliance, its default state as a network is Unmanaged.

In the **Assets** category:

- **Discovered:** The total number of Assets discovered by Probe appliances.
- **Managed:** The number of discovered assets that are currently managed by the NIOS Grid. These devices have been converted from Unmanaged status to Managed status.
- **Unmanaged:** The number of IPs with discovered data that are counted as Unmanaged by the NIOS Grid Master, and have not been converted into a Host or a Fixed IP Address. After an Asset is discovered and catalogued by a Probe appliance, its default state is Unmanaged.
- **Conflicts:** The number of discovered assets that have a MAC address conflict or are part of a configured DHCP range, but do not have a fixed address or lease record and are not part of an exclusion range.

## My Commands

The *My Commands* widget provides easy access to commands that you frequently use, so you can perform your tasks without leaving the Dashboard. You can add one *My Commands* widget to your Dashboard.

To configure the *My Commands* widget, click the Configure icon and do the following:

- Select a command from the **Available** list and click the > arrow to move it to the **Selected** list. You can always toggle the commands between the two lists. Select multiple commands by using SHIFT-click and CTRL-click.

## DDNS Statistics

The *DDNS Statistics* widget provides information about the dynamic DNS (DDNS) updates that occur on the DNS service of a selected Grid member. The widget displays the total number of DDNS updates that succeeded, failed, and that were rejected. It also displays a line graph that tracks the status of the DDNS updates per second.

You can add a *DDNS Statistics* widget to your Dashboard for each DNS server on the Grid that accepts dynamic DNS updates.

To configure the *DDNS Statistics* widget, click the Configure icon and do the following:

- Click **Select Member**. In the *Member Selector* dialog box, select a Grid member from the list.

- **Graph Configuration:** Select which updates you want to track in the **Updates per Second** graph:
  - **Success:** The number of DDNS update requests that succeeded.
  - **Prerequisite Reject:** The number of DDNS update requests that were rejected because the prerequisite conditions specified in the request were not met.
  - **Reject:** The number of DDNS update requests that were rejected by the DNS service.
  - **Failure:** The number of DDNS update requests that failed.

The widget displays the following information:

- **DDNS Updates** tab: Displays totals for each type of update.
- **Updates per Second** tab: Displays a line graph that tracks the status of the DDNS updates. The time is displayed according to the time zone specified in the User Profile. If the auto-detect time zone option is enabled and Grid Manager cannot determine the browser time zone, then the time is displayed in UTC format. You can mouse over the graph to display the coordinates of any point in the graph.

## System Activity Monitor

The *System Activity Monitor* widget provides information about the following resources on the selected Grid member: CPU and its utilization, system memory, NIC usage, top processes, and information about VLAN interfaces. By default, the widget displays the system activity of the Grid Master. You can add a *System Activity Monitor* widget to your Dashboard for each Grid member whose resources you want to monitor.

To configure the *System Activity Monitor* widget, click the Configure icon and select a Grid member and the resources that you want to track:

- Click **Select Member**. In the *Member Selector* dialog box and select a Grid member from the list.
- **CPU:** Select which type of CPU usage you want to track:
  - **User:** The CPU usage of user applications, such as programs and libraries.
  - **System:** The CPU usage of the kernel and drivers.
  - **Idle:** The percentage of CPU that is not in use.
- **System Memory:** Select which portion of the system memory you want to track:
  - **Real Memory Used:** The physical RAM usage.
  - **Swap Used:** The swap area usage. The swap area is the disk area that temporarily holds a process memory image.
- **NIC Usage:** Select how you want to measure network traffic:
  - **Bytes:** Reports the number of bytes.
  - **Packets:** Reports the number of packets.
- **NIC Settings:** Select the port on which you want to measure network traffic. If you have configured VLANs, Grid Manager displays them in the format LAN1 nnnn or LAN2 nnnn, where nnnn represents the associated VLAN ID. For example, a VLAN configured on LAN1 can be displayed as LAN1 297 and a LAN2 VLAN can be LAN2 21. For more information about VLANs, see [VLAN Management](#).



### Note

For vNIOS appliances, some of the options in the drop-down list may vary depending on your vNIOS configuration. For example, if you are using a single network interface instance of vNIOS for GCP, you will see choices specific to the LAN1 interface only. For more information, see the vNIOS documentation specific to your product at [Appliances](#).

- **CPU Utilization and Top N Processes:** Set the auto refresh period in this section. NIOS displays the information for all available cores.
  - **Auto Refresh Period for CPU Utilization and Top N Processes:** Enter the time interval in seconds for the CPU Core Utilization graph and the top N process data to auto refresh and display the CPU core utilization information. If you enter 12, the graph displays new information after every 12 seconds. You can enter a minimum refresh interval of 10 seconds and a maximum refresh interval of 30 seconds. By default, the time interval is set to 10 seconds. This field is applicable only to the **CPU Utilization** and **Top N Processes** tabs.
- **Auto Refresh Period:** Enter the refresh interval in seconds for the data in the **CPU**, **System Memory**, and **NIC Usage** tabs to auto refresh.

The *System Activity Monitor* widget displays a tab for each resource: **CPU**, **System Memory**, **NIC Usage**, **CPU Utilization**, **Top N Processes**.

Each tab contains a line graph that tracks the resource utilization per second.

- **CPU:** The graph on the **CPU** tab tracks the percentage of CPU usage.
- **System Memory:** The graph on the **System Memory** tab tracks the memory utilization percentage.
- **NIC Usage:** The graph on the **NIC Usage** tab tracks either bytes or packets per second.
- **CPU Utilization:** If you select the **Live** option, the graph tracks live CPU utilization data for the last 10 minutes for all CPUs in your Grid member. The graph is refreshed based on the time interval you specify in the **Auto Refresh Period for CPU Utilization and Top N Processes** field. Each CPU is denoted in a different color. If you select the **Historical** option, you can view the CPU utilization data for up to a maximum of past 60 minutes based on the time range you specify in the **Earliest** and **Latest** fields. For example, if you enter 2019-09-05 and 09:20:42 AM in the **Earliest** field and 2019-09-05 and 10:20:42 AM in the **Latest** field, the graph displays the CPU utilization data for 5th September 2019 between 9:20:42 AM and 10:20:42 AM. You can view data for a maximum of past of 5 days but the time difference between **Earliest** and **Latest** time should not exceed 60 minutes.
- **Top N Processes:** If you select the **Live** option, the table displays the process ID and name of the top N processes that are consuming CPU utilization. N is the number that you specify in the **Number of Top Processes** field on the **Monitoring** tab of the *Grid Properties* editor. It also displays the percentage of CPU utilized by each process. The data is refreshed based on the time interval you specify in the **Auto Refresh Period for CPU Utilization and Top N Processes** field. If you select the **Historical** option, you can view past top N process data based on the time range you specify in the **Earliest** and **Latest** fields. For example, if you enter 2019-09-05 and 09:20:42 AM in the **Earliest** field and 2019-09-05 and 10:20:42 AM in the **Latest** field, the graph displays the top process data on 5th September 2019 between 9:20:42 AM and 10:20:42 AM. You can view data for a maximum of 5 days.

The time is displayed according to the time zone specified in the User Profile. If the auto-detect time zone option is enabled and Grid Manager cannot determine the browser time zone, then the time is displayed in UTC format. You can mouse over the graph to display the coordinates of any point in the graph.

## File Distribution Statistics

The *File Distribution Statistics* widget enables you to monitor the status of file distributions services from the Dashboard. The widget provides an overall status of file distribution on all members in the Grid. It also displays the file system utilization for the file distribution subsystem.

The service status displays one of the following:

- **OK:** All file distribution services are running properly.
- **Stopped:** All file distribution services are stopped.
- **Warning:** The file distribution services are not running properly.
- **Error:** The file distribution services encounter an error.

You can click the link to view detailed information about the file distribution services. Grid Manager displays the Members tab on the File Distribution tab.

To configure the *File Distribution Statistics* widget, click the Configure icon and select one of the following chart types:

- **Pie**
- **Bar**

The *File Distribution Statistics* widget displays the following information:

- **File System Utilization:** The percentage of utilization of the overall allocated file distribution subsystem space on all members. You can use this information to verify if there is sufficient space for file distribution in the Grid.

## Active WebUI Users

The *Active WebUI Users* widget provides information about the users who are logged in to Grid Manager or System Manager. It does not include users who are using the Infoblox API or are logged in to the serial console.

You can add only one *Active WebUI Users* widget to the Dashboard. You must have a superuser account to add this

widget to the Dashboard.

It displays the following information about each user:

- **User ID:** The user name.
- **Source Address:** The IP address of the management station the user used to connect to Grid Manager.
- **Logged In Since:** The date and time the user logged in.
- **Idle Time:** The number of minutes the user has not had any activity on Grid Manager. Note that the idle session timeout is 2 hours, so the idle time is cleared every 2 hours.
- **User Agent:** The system used to access Grid Manager, such as the browser version and platform information. You can sort the columns and hide or display each one. You can also export the list to a .csv file.





### Microsoft Servers Status Widget

The *Microsoft Servers Status* widget displays the operational status of each Microsoft server managed by the Grid. Grid Manager displays this widget only when at least one member in the Grid has a Microsoft management license. You can configure this widget to display the status of all Microsoft servers or only those with warnings and errors. You can also view the monitor and control status for the DNS and DHCP service on the Microsoft server. To modify the *Microsoft Servers Status* widget, click the Configure icon and select one of the following:





- **Show all Microsoft servers**
- **Only show servers with service warnings or errors**

The *Microsoft Servers Status* widget displays the following information about each Microsoft server:

- **Server Name:** The hostname of the Microsoft server.
- **IP Address:** The IP address of the Microsoft server.
- **Status:** The connection status of the Microsoft server.
  - **OK:** The Grid member is connected to the Microsoft server.
  - **Connecting:** The Grid member is connecting to the Microsoft server.
  - **Error:** The Grid member tried to connect to the Microsoft server, but failed. You can check the syslog for any messages.
  - **Not Available:** The Microsoft server is disabled. The Grid member does not try to connect to disabled servers.
- **DNS:** The status of the DNS service on the Microsoft server. The DNS service status can be one of the following:

Icon	Color	Meaning
	Green	The DNS service is functioning properly.
	Red	The Microsoft server is unavailable.
	Yellow	The DNS service is starting or stopping.
	Gray	The DNS service is stopped or management of the Microsoft DNS server is disabled.

- **DHCP:** The status of the DHCP service on the Microsoft server. The DHCP service status can be one of the following:

Icon	Color	Meaning
	Green	The DHCP service is functioning properly.
	Red	The Microsoft server is unavailable.
	Yellow	The DHCP service is starting or stopping.
	Gray	The DHCP service is stopped or management of the Microsoft DHCP server is disabled.

- **Active Directory Sites:** The status icon in green indicates the synchronization status of Active Directory Sites on the Microsoft server.
- **Enable DNS Monitor & Control:** Displays **Yes** if the monitor and control status is enabled for the DNS service on the Microsoft server and displays **No** if it is disabled.
- **Enable DHCP Monitor & Control:** Displays **Yes** if the monitor and control status is enabled for the DHCP service on the Microsoft server and displays **No** if it is disabled.

## CSV Import Manager

The **CSV Import Manager** on the Status Dashboard displays the status of CSV import jobs you have submitted. You can start a file import from **CSV Import Manager** and control and monitor it from this widget. You can also launch **CSV Import Manager** from the Task Dashboard or the Toolbar. You can also delete uploaded CSV files. For more information, see [Importing and Exporting Data using CSV Import](#). You can click the Refresh icon or configure auto refresh to update the status.

The widget displays the following information about the import jobs that were submitted in the past 30 days:

- **User Name:** The admin user who submitted the CSV import. Only superusers can view this column.
- **Status:** The current status of the import job. The status can be one of the following:
  - **Import successful:** The import is completed without errors. Check the **Message** field for information about the import.
  - **Import unsuccessful:** The import is completed, but with errors. Check the **Message** field for information about the error message.
  - **Import pending:** The job is in queue for execution.
  - **Import in progress:** The job is being executed.
  - **Import stopped:** The job has been stopped. You can select the job and restart the import.
  - **Test successful:** Test is completed without errors. Check the **Message** field for information about the test.
  - **Test unsuccessful:** Test is completed, but with errors. Check the **Message** field for information about the error message.
  - **Test pending:** Test is in queue for execution.
  - **Test in progress:** Test is in progress.
  - **Test stopped:** Test has been stopped. You can select the job and restart the import.
  - **Saved file:** The data file has been uploaded, but the import has not started. Note that after a product restart, which can be caused by a failover, all **Import in progress** jobs go into **Import stopped** state; all **Import pending** jobs continue to be queued for execution.
  - **Submitted:** The timestamp when the job was submitted.
  - **Completed:** The timestamp when the job was completed. This field is blank if the job has not been completed yet.
  - **File Name:** The CSV data file name.
  - **Message:** This field displays the number of rows of data that has been processed and the number of rows of data the import has detected errors. Depending on the import options, Grid Manager displays the row

number at which it stops the import when it encounters an error, or the total number of rows it has processed by skipping over the erroneous data. For example, if there are 100 rows of data and you select "On error: Stop importing," and there is an error in row 90, the appliance displays **90of100completed,1error**. If you select "On error: Skip to the next row and continue," the appliance displays **100 of 100 completed, 1 error**.

- **File Size:** The CSV data file size.

#### Note

Superusers can view all CSV import jobs and limited-access users can view only the jobs they submitted.

## Pending Approvals

The *Pending Approvals* widget provides information about tasks that are pending your approvals. Add the *Pending Approvals* widget to your Dashboard to monitor tasks that require your approvals.

You can select a task and perform the following:

- Click the Approve icon to approve the task.
- Click the Reject icon to disapprove the task.

You can also click **Task Manager** to access the **Administration** tab -> **Workflow** tab -> **Task Manager** tab.

The *Pending Approvals* widget displays the following information about each task that requires your approval:

- **Task ID:** The ID associated with the task. The appliance assigns an ID to a task in chronological order.
- **Submitter:** The username of the admin who scheduled or submitted the task.
- **Ticket Number:** The reference number entered by the submitter to identify the task. You can enter up to 20 alphanumeric characters.
- **Scheduled Time:** The date, time, and time zone when the task was scheduled for execution.
- **Affected Object:** The name or value of the object that is associated with the task. For example, if the task involves an A record, this field displays the domain name of the record. If it is a fixed address, it displays the IP address of the fixed address.
- **Object Type:** The object type. For example, the appliance can display A Record or Fixed Address.
- **Action:** The operation the appliance performs in this task. The operation can be: **Add**, **Modify**, **Delete**, or **Network Discovery**.
- **Submitte Time:** The date, time, and time zone when the task was submitted. You can select this for display. It is not displayed by default.

## Infoblox Community

The *Infoblox Community* widget displays the latest news from Infoblox. It provides links to video clips that show you how to perform certain tasks, such as how to prepare for IPAM Express and how to add a network. You can click available links in the widget to get more information about Infoblox products and solutions.

Note that content in the *Infoblox Community* widget may not be displayed in certain versions of Mozilla FireFox, Google Chrome, and Microsoft Internet Explorer due to restrictions these browsers use to block certain secure data.

Follow these steps to unblock the *Infoblox Community* widget and view data in your respective browser:

- **MozillaFireFox:** Click the *Shield* icon in the address bar and choose **DisableProtectiononThisPage** from the drop-down list. The icon in the address bar changes to a warning triangle and content is displayed in the *InfobloxCommunity* widget. For more details, refer to information at <https://blog.mozilla.org/tanvi/2013/04/10/mixed-content-blocking-enabled-in-firefox-23/>.
- **Google Chrome:** Click the *Shield* icon in the address bar and click **Load unsafe script** in the pop-up box. Chrome automatically refreshes the webpage and loads the content in the *Infoblox Community* widget. For more details, refer to information at <https://support.google.com/chrome/answer/1342714?hl=en>.
- **Internet Explorer:** Click the *Compatibility View* icon adjacent to the address bar. The browser refreshes and the *Security Warning* dialog box is displayed. Click **No** in the dialog box. The **Only Secure content is displayed** pop-up blocker is displayed at the bottom of the browser. Click the **Show all content** button in this pop-up blocker to view the content. For more details, refer to the information at <http://windows.microsoft.com/en-in/internet-explorer/use-compatibility-view#ie=ie-8>.



## Mobile Devices Status

The *Mobile Devices* widget provides information about the number of active leases of the DHCP fingerprint devices managed by the Grid. The widget displays a pie chart indicating the number of active leases in percentile for each of the device category. You can click the Refresh icon or configure auto refresh to update the status.



### Note

The *Mobile Devices* widget updates its data every 15 minutes. A device might not be displayed in this widget if its lease expires within 15 minutes.

To configure the *Mobile Devices* widget, click the Configure icon and do the following:

- Click **Select Network View**. In the *Network View Selector* dialog box, select a network view from the list and click **OK**.

Note that if multiple network views were previously configured, Grid Manager displays the default network view. You can select another network view from the *Network View Selector* dialog box.

The widget displays the number of active leases for the following device classes:

- **MacOS** - Displays all devices that were detected to be running Mac OS.
- **Windows** - Displays all devices that were detected to be running Windows.
- **Android Mobile** - Displays Smartphones/PDAs/Tablets that were detected to be running Android.
- **Apple Mobile** - Displays Smartphones/PDAs/Tablets that were detected to have Apple in the DHCP fingerprint information.
- **No Match** - Displays all devices whose fingerprint information does not match with any of the standard/custom DHCP fingerprint data stored in the appliance. For information about Standard and Custom DHCP Fingerprints, see [Standard and Custom DHCP Fingerprints](#).
- **Other** - Displays all devices that belong to a device class other than those listed above.

### List of device types and classes

Category	Device Class	Device Type
Windows	Windows	Microsoft Windows 2000
		Microsoft Windows 2003
		Microsoft Windows 8
		Microsoft Windows Vista/7 or Server 2008
		Microsoft Windows XP
Mac OS	Macintosh	Apple Mac OS 9
		Apple Mac OS X, TV (HD)
Apple Mobile	Smartphones/PDAs/Tablets	Apple iPod
		Apple iPod, iPhone, iPad or TV (SD)



Category	Device Class	Device Type
Android Mobile	Smartphones/PDAs/Tablets	Android Phone/Tablet (Generic)
		Android Phone/Tablet (HTC, older devices)
		Android Phone/Tablet (Motorola, older devices)
		Android Phone/Tablet (Sony Ericsson, older devices)
		Android Phone/Tablet (Unknown devices)
		Android Phone/Tablet (Vizio tablet, Others)
		Android Phone/Tablet (newer devices)
		Android tablets (Samsung, Others)
		ZTE N9120 Android

## DNS Integrity Check

The DNS Integrity Check widget displays status about DNS data discrepancies that have been detected through DNS integrity check that is designed to mitigate DNS domain hijacking. This widget displays top-level or parent authoritative zones that have been selected for DNS data monitoring. For information about how to configure DNS integrity check to mitigate possible DNS domain hijacking, see [Configuring DNS Integrity Check for Authoritative Zones](#).

The widget displays the following information (note that this table is sorted by **Status**):

- Left click the Action icon next to a zone to perform the following:
  - **View Syslog**: Select this to open the *Syslog Preview* dialog and view data discrepancy events for the selected zone.
  - **Check Now**: Select this to perform DNS integrity check to immediately query current DNS data from the top-level parent domain. When you select this, verbose logging for DNS integrity check is automatically enabled. After the operation is complete, the appliance updates the timestamp for the **Last Checked** column.
- **Zone**: Displays the name of the top-level authoritative zones that is being monitored for DNS integrity check. You can click the zone name and the appliance opens the zone viewer for the selected zone.
- **Status**: Displays the current DNS data discrepancy status. The status can be one of the following:
  - **Critical** (red): Data in the NS RRsets for the authoritative and delegate zones are completely out of synchronization.
  - **Severe** (orange): Some data in the NS RRset between the authoritative and delegate zones overlaps and some data is different.
  - **Warning** (yellow): The NS RRset for the authoritative zone is a subset of the NS RRset for the delegate zone. It is possible that incorrect IP addresses have been entered at the registrar.
  - **Informational** (blue): The NS RRset for the delegate zone is a subset of the NS RRset for the authoritative zone. This could indicate a possible delay in domain registration.
  - **Normal** (green): There are no DNS data discrepancies between the NS RRsets for the authoritative and delegated zones.
  - **None** (black): No DNS discrepancies data has been collected or DNS integrity check has not been performed.
- **Last Checked**: The timestamp in YYYY-MM-DD HH:MM:SS when the parent domain was last queried for its DNS data.

- **Description:** Information about the zone.

## Previewing Syslog Events

When you select **View Syslog** from the *DNS Integrity Check* widget for a selected zone, the *Syslog Preview* dialog is displayed. You can view related syslog events for the selected zone in this dialog, as follows:

- **Timestamp:** The timestamp in YYYY-MM-DD HH:MM:SS when the event was logged.
- **Facility:** The location that determines the processes and daemons from which the log messages are generated.
- **Level:** The severity level of the DNS data discrepancies. This can be **Critical**, **Severe**, **Warning**, **Information**, or **Normal**. For more information, see descriptions for the **Status** field.
- **Server:** The name of the Grid member that performed the data check.
- **Message:** Syslog information about the event. If you have enabled verbose logging, this displays detailed information about the event. For information about how to enable verbose logging, see [Configuring DNS Integrity Check for Authoritative Zones](#).

You can also click **Go to Syslog Viewer** on the upper right corner of the dialog to view all events in the syslog. For more information about the syslog, see [Viewing the Syslog](#).

## Cloud Statistics

The *Cloud Statistics* widget appears only when you have deployed the Cloud Network Automations license on the Grid Master. This widget displays statistical information for cloud objects. It contains the following tabs: *Tenant & VMs*, *Fixed vs. Floating* and *Available vs. Allocated*. You must install valid cloud related licenses to access this widget. For more information about installing licenses and enabling Cloud Network Automation, see [Deploying Cloud Network Automation](#).

To modify the *Cloud Statistics* widget, click the Configure icon and select one of the following:

- **Show Statistics From:**
  - **All Tenants:** Select this to display statistics for all tenants.
  - **Select Tenant:** Click **Select** to choose a specific tenant for which statistics are displayed.
- **Show:**
  - **All IP Addresses:** Select this to display all IP address allocation for all tenants or the tenant of your choice.
  - **Fixed:** Select this to display only fixed IP address allocation for all tenants or the tenant of your choice. Fixed IP addresses correspond to OpenStack Fixed IP Addresses.
  - **Floating:** Select this to display only floating IP address allocation for all tenants or the tenant of your choice. Floating IP addresses correspond to OpenStack Floating IP Addresses.

## Dig Request

The *Dig Request* widget enables you to perform a DNS lookup on the Grid Master or on the specified Grid member and displays the output of the dig command.

### Note

When RPZ license is installed on both the Grid Master and the Grid member, the RPZ rule might not be triggered if you perform dig on the Grid member from the Grid Master.

To perform a DNS lookup using the dig command, complete the following:

- **Run dig command on:** Select one of the following. The default is **Grid Master**.
  - **Grid Master:** Select this to perform a DNS lookup on the Grid Master.
  - **Grid Member:** Select this to perform a DNS lookup on the Grid member. Click **Select Member** to select a Grid member. If there are multiple members, the *Member Selector* dialog box is displayed, from which you can select a member. Click the required member name in the dialog box. You can also click **Clear** to clear the displayed member and select a new one.
- **Name Server to Query(Optional):** Optionally, specify the name server on which you want to perform a DNS lookup. You can enter either the name, IPv4 address, or IPv6 address of the name server.

- **Record Type:** Select the resource record type from the drop-down list. You can select **Any** to query all the resource record types or select one of the following from the drop-down list: **A, AAAA, CAA, CNAME, DNAME, MX, NAPTR, NS, PTR, SRV, TXT, AXFR**. If the record type is **Unknown**, then directly enter the type of unknown record. For example, for an unknown record of type RP, enter **RP**. The default is **Any**.
- **Send Recursive Query:** Select this to send recursive queries for the domain. This checkbox is selected by default.
- **Domain Name to Query:** Enter the domain name to query.

Click **Perform Dig**. The widget displays the status and output of the dig command.

Note that if you have installed RPZ license and enabled RPZ logging in the Grid, you can view RPZ syslog messages by clicking **View RPZ Syslog** if the specified domain name matches the RPZ rule.

## Security Status for Grid

The *Security Status for Grid* widget displays the overall status of Threat Protection, RPZ (Response Policy Zone), and DNS Threat Analytics services on the Grid members that support Infoblox Advanced DNS Protection, hardware or Software ADP, and Infoblox Threat Insight. Grid Manager displays this widget only when at least one member in the Grid has the Threat Protection, RPZ, or Threat Analytics license installed. You can add this widget to the Security dashboard to monitor the overall security status of the Grid. The statistics displayed in this widget are cumulative, collected from all the Grid members that support Infoblox Advanced DNS Protection, hardware or Software ADP, and Infoblox Threat Insight. This widget displays data for the last 30 minutes. The overall status of Threat Protection, RPZ, and DNS Threat Analytics is determined by the threshold values configured in the *Global Dashboard Properties* editor. For information, see *Configuring Security Status Thresholds* below.

### Note

If the Threat Protection license is not installed on any of the Grid members, Grid Manager does not display any threat protection related information in this widget. Similarly, if the RPZ license is not installed on any of the Grid members, Grid Manager does not display RPZ and DNS Threat Analytics related information in this widget and if the Threat Analytics license is not installed on any of the Grid members, Grid Manager does not display DNS Threat Analytics related information in this widget.

The widget displays the following information for Threat Protection, RPZ, and DNS Threat Analytics:

- **Status:** It displays the overall status of the security service in the Grid based on the events collected from all the members that support Infoblox Advanced DNS Protection and Infoblox Threat Insight. It represents the status of the most critical member in the Grid.

The status icon can be one of the following for the Threat Protection, RPZ, and DNS Threat Analytics service:

- **OK (Green):** The license for the security service is installed and the security service is running. The rulesets for the security service are available and the number of events triggered are less than the yellow and red threshold values configured in the *Global Dashboard Properties* editor for the corresponding security service.
- **Warning (Yellow):** The license for the security service is installed and the security service is running. The rulesets for the security service are available and the number of events triggered for any of the parameters equals or exceeds the yellow threshold value, but less than the red threshold value configured in the *Global Dashboard Properties* editor for the corresponding security service.
- **Critical (Red):** The license for the security service is installed and the security service is running. The rulesets might not be available or the number of events triggered for any of the parameters, equals or exceeds the red threshold value configured in the *Global Dashboard Properties* editor for the corresponding security service.
- **Not Setup (Black):** The license for the security service is installed, but the security service is not running.
- **Unknown (Black):** The data is not available from the Grid member.

You can hover your mouse over the Threat Protection, RPZ, and Threat Analytics status icon and view the *Threat Protection Status for Grid* widget, *Response Policy Zone (RPZ) Status for Grid* widget, and *Threat Analytics Status for Grid* widget respectively. For information about *Threat Protection Status for Grid* widget, *Response Policy Zone (RPZ) Status for Grid* widget, and *Threat Analytics Status for Grid* widget, see below *Threat Protection Status for Grid*, *Response Policy Zone (RPZ) Status for Grid*, and *Threat Analytics Status for Grid* respectively.

- **Events from <> of <> security capable members:** This column displays the cumulative event counts collected from the online Grid members that support the Infoblox Advanced DNS Protection and Infoblox Threat Insight.

- **Threat Protection:** Displays the total threat protection event counts for the following severity levels:
  - **Critical** (Red): The total number of critical events.
  - **Major** (Orange): The total number of major events.
  - **Warning** (Yellow): The total number of warning events.
  - **Informational** (Blue): The total number of informational events.
- **RPZ:** Displays the total number of hits received for the following RPZ rules:
  - **Blocked hits** (Red): Total number of queries that triggered a Block (No Data) or Block (No Such Domain) RPZ rule.
  - **Passthru hits** (Yellow): Total number of queries that triggered a Passthru RPZ rule.
  - **Substituted hits** (Orange): Total number of queries that triggered a Substitute (Domain Name) or Substitute (Record) RPZ rule.
- **Analytics:** Displays the total number of DNS tunneling events.
- **Definitions/Rules:** This column displays the status of the latest ruleset available in the database. For RPZ, the definition status is based on the latest RPZ feed received from Infoblox specific feeds. You can hover your mouse over the definition status to see the RPZ definition status when RPZ definitions exists.
- **Configuration Status:** This column indicates whether the security service is enabled and running properly or not. Grid manager displays a green check mark if the security service is enabled and running properly in the Grid. If the security service is disabled, a gray pause mark is displayed. You can hover your mouse over the gray pause mark to see the status of the security service.

In addition, you can click the Configure icon and do the following:

- Click **Configure Security Status Thresholds** to configure the thresholds for the security status of the Grid. In the *Global Dashboard Properties* editor, you can define the threshold values for Threat Protection, RPZ, and DNS Threat Analytics. For information, see [Configuring Security Status Thresholds](#) below.
- Select the **Auto Refresh Period** checkbox to turn on auto-refresh and specify the auto-refresh period in seconds. The default auto-refresh period is 30 seconds.

#### **Warning**

If the [Detailed Status panel](#) is open, the following actions take place:

- Grid Manager auto refreshes at a rate of 30 seconds.
- Widgets that support user-specified auto refresh, refresh at the rate specified in the **Auto Refresh Period** field.

Therefore, even if you set the session timeout to be to greater than the auto refresh rate, auto refresh still takes place. The Grid Manager session does not time out because auto refresh takes precedence over the session timeout.

Click the Configure icon again to hide the configuration panel after you complete the modification.

## Security Status for All Members

The *Security Status for All Members* widget displays information about the status of all the Grid members that support Infoblox Advanced DNS Protection, hardware or Software ADP, and Infoblox Threat Insight. Grid Manager displays this widget only when at least one member in the Grid has the Threat Protection or RPZ license. You can add this widget to the Security dashboard to monitor the status of the Grid members that support Infoblox Advanced DNS Protection, hardware or Software ADP, and Infoblox Threat Insight.

#### **Note**

When an HA Grid Master fails over, the new active node re-collects data from all the Grid members. Hence, it might take a few seconds until the data is displayed in the Security dashboard. When an HA Grid member fails over, the Grid Master stops collecting data from the HA member.

The *Security Status for All Members* widget displays the following information:

- **Overall Status:** The current overall security status of the members that support Infoblox Advanced DNS Protection and Infoblox Threat Insight. This can be **OK**, **Warning**, **Critical**, or **Unknown**.

The security status for a member might be **Unknown** if NTP service is out of synchronization for the member. Hence, to ensure that the correct data is displayed for the member, use NTP to synchronize the time of the member with that of the Grid Master.

- **Member:** The name of the member. You can hover your mouse over the member name and view the *Member Status* widget. For information about the *Member Status* widget, see Member Status (System Status) below.
- **IPv4 Address:** The IPv4 address of the member.
- **IPv6 Address:** The IPv6 address of the member.
- **Threat Protection Status:** The status of the threat protection service running on the member. This can be either **OK**, **Warning**, **Critical**, **NotSetup**, or **Unknown**. You can hover your mouse over the threat protection status and view the *Threat Protection Status for Member* widget. For information about the *Threat Protection Status for Member* widget, see Threat Protection Status for Member below.
- **RPZ Status:** The status of the RPZ service running on the member. This can be either **OK**, **Warning**, **Critical**, **NotSetup**, or **Unknown**. You can hover your mouse over the RPZ status and view the *ResponsePolicyZone(RPZ)Statistics* widget. For information about the *Response Policy Zone (RPZ) Statistics* widget, see Response Policy Zone (RPZ) Status for Member below.
- **Analytics Status:** The status of the DNS Threat Analytics service running on the member. This can be either **OK**, **Warning**, **Critical**, **NotSetup**, or **Unknown**.

You can also do the following in this widget:



#### Warning

If the [Detailed Status panel](#) is open, the following actions take place:

- Grid Manager auto refreshes at a rate of 30 seconds.
- Widgets that support user-specified auto refresh, refresh at the rate specified in the **Auto Refresh Period** field.

Therefore, even if you set the session timeout to be to greater than the auto refresh rate, auto refresh still takes place. The Grid Manager session does not time out because auto refresh takes precedence over the session timeout. For more information about widgets, see [Status Dashboard](#).

- Turn on auto-refresh. Click the Configure icon and select the **AutoRefreshPeriod** checkbox to turn on auto-refresh. Specify the auto-refresh period in seconds. The default auto refresh period is 30 seconds.
- Click the Action icon (shown as a gear in each row of the table) next to the overall status of each member, and select **ViewSyslog** to view all the events logged in the syslog. Grid Manager displays the syslog messages in the Syslog Preview window.
- Click the Export icon to export the data displayed in this widget.
- Click the Print icon to print the data displayed in this widget.
- Click **Response Policy Zones** link in the **Go To** field at the top of the widget to view the RPZs configured on the member. Grid Manager displays the **Response Policy Zones** tab on the **DNS** tab. To navigate back to the Security dashboard, click **Back to Security Dashboard** at the top left corner of the navigation bar in the **Response Policy Zones** tab.
- Click **Threat Protection** link in the **Go To** field at the top of the widget to view the threat protection rulesets configured on the member. Grid Manager displays the **Threat Protection Rules** tab on the **Security** tab. To navigate back to the Security dashboard, click **Back to Security Dashboard** at the top left corner of the navigation bar on the **Threat Protection Rules** tab.
- Click **Threat Analytics** link in the **Go To** field at the top of the widget to view the whitelist domains configured on the member. Grid Manager displays the **Threat Analytics** tab. To navigate back to the **Security** dashboard, click **Back to Security Dashboard** at the top right corner of the panel on the **Threat Analytics** tab.
- Click **Members** link in the **Go To** field at the top of the widget to view the members configured in the Grid. Grid Manager displays the **Members** tab on the **Grid Manager** tab. To navigate back to the Security dashboard, click **Back to Security Dashboard** at the top left corner of the navigation bar on the **Members** tab.

## Threat Protection Status for Grid

The *Threat Protection Status for Grid* widget displays the statistical information about the threat protection events triggered on all the members in the Grid that support Infoblox Advanced DNS Protection, hardware or Software ADP, and Infoblox Threat Insight. This widget contains the following tabs: *Total Events by Severity*, *Top 10 Grid Members*, *Events*

*Over Time, Top 10 Rules, and Top 10 Clients.*  
You can do the following in this widget:

 **Warning**

If the *Detailed Status panel* is open, the following actions take place:

- Grid Manager auto refreshes at a rate of 30 seconds.
- Widgets that support user-specified auto refresh, refresh at the rate specified in the **Auto Refresh Period** field.

Therefore, even if you set the session timeout to be to greater than the auto refresh rate, auto refresh still takes place. The Grid Manager session does not time out because auto refresh takes precedence over the session timeout.

- Turn on auto-refresh.  
Click the Configure icon, select the **Auto Refresh Period** checkbox, and specify the refresh period in seconds. The default auto refresh period is 30 seconds. You can click the Configure icon again to hide the configuration panel.
- Click the **Total Events by Severity** tab to view information about threat protection related events by the severity level.
- Click the **Top 10 Grid Members** tab to view information about the top 10 Grid members that have the most number of threat protection events.
- Click the **Events Over Time** tab to view information about the total event count for each type of event severity in the given time frame.
- Click the **Top 10 Rules** tab to view information about the top 10 threat protection rules with the most number of hits.
- Click the **Top 10 Clients** tab to view information about the top 10 clients that have the most number of threat protections events.

### Total Events by Severity

The **Total Events by Severity** tab displays statistics about the Threat Protection events for each type of event severity. This tab displays a bar chart that lists the total event counts for each severity level. Each severity level is represented by a different color. The event statistics are cumulative, collected from all the members in the Grid that support Infoblox Advanced DNS Protection and Infoblox Threat Insight.

This line graph displays the event counts for the following severity levels:

- **Critical** (Red): The total number of critical events.
- **Major** (Orange): The total number of major events.
- **Warning** (Yellow): The total number of warning events.
- **Informational**(Blue): The total number of informational events.

If you have configured a reporting member in the Grid, the **Go To History** link is displayed on this tab. You can click **Go To History** to view the *Threat Protection Event Count By Severity Trend* report on the **Reporting** tab. To navigate back to the Security dashboard from the **Reporting** tab, click **Back to Security Dashboard** at the top left corner of the navigation bar in the **Reporting** tab.

### Top 10 Grid Members

The **Top 10 Grid Members** tab displays a stacked bar chart that tracks the top Grid members with the most total counts of threat protection events. Each severity level is represented with a different color. The report displays the top 10 members in descending order.

If you have configured a reporting member in the Grid, the **Go To History** link is displayed on this tab. You can click **Go To History** to view the *Threat Protection Event Count By Member/Member Group Trend* report on the **Reporting** tab. To navigate back to the Security dashboard from the **Reporting** tab, click **Back to Security Dashboard** at the top left corner of the navigation bar on the **Reporting** tab.



## Events Over Time

The **Events Over Time** tab displays a line graph that tracks the event count for each event severity in a given time frame. You can view the event counts for the following severity level: Critical, Major, Warning, and Informational. The event statistics are cumulative, collected from all the members in the Grid that supports Infoblox Advanced DNS Protection and Infoblox Threat Insight. Each severity level is represented with a different color.

If you have configured a reporting member in the Grid, the **Go To History** link is displayed on this tab. You can click **Go To History** to view the *Threat Protection Event Count By Severity Trend* report on the **Reporting** tab. To navigate back to the Security dashboard from the **Reporting** tab, click **Back to Security Dashboard** at the top left corner of the navigation bar on the **Reporting** tab.

## Top 10 Rules

The **Top 10 Rules** tab displays a horizontal bar chart that tracks the top threat protection rules that have the most number of hits. Each severity level is represented with a different color. The report displays the top 10 rules in descending order. If you have configured a reporting member in the Grid, the **Go To History** link is displayed in this tab. You can click **Go To History** to view the *Threat Protection Top Rules Logged* report in the **Reporting** tab. To navigate back to the Security dashboard from the **Reporting** tab, click **Back to Security Dashboard** at the top left corner of the navigation bar on the **Reporting** tab.

## Top 10 Clients

The **Top 10 Clients** tab displays a horizontal bar chart that tracks the total number of threat protections events triggered by top clients (source IP addresses). This tab displays the IP addresses of the top 10 clients. For NAT clients, it displays the NAT addresses for the clients.

If you have configured a Reporting member in the Grid, the **Go To History** link is displayed on this tab. You can click **Go To History** to view the *Threat Protection Top Rules Logged by Source* report on the **Reporting** tab. To navigate back to the Security dashboard from the **Reporting** tab, click **Back to Security Dashboard** at the top left corner of the navigation bar on the **Reporting** tab.

### Note

The data displayed in this widget may not be consistent with the data displayed in the *Threat Protection Top Rules Logged by Source* report.

## Threat Protection Status for Member

The *Threat Protection Status for Member* widget displays statistics about the threat protection events for a specific Grid member that supports Infoblox Advanced DNS Protection, hardware or Software ADP. For information about the threat protection feature, see [About Infoblox Advanced DNS Protection](#).

To configure the *Threat Protection Status for Member* widget, click the Configure icon and complete the following:

- Click **Select Member**. In the *Member Selector* dialog box, select a Grid member from the list that supports Infoblox Advanced DNS Protection, hardware or Software ADP.
- Select either **Dial** or **Bar** as the display format for the following resources: **Smart NIC CPU**, **Traffic being dropped**, **Traffic being received**. Note that Smart NIC CPU selection is displayed only when you select a Grid member that supports Infoblox Advanced DNS Protection.
- **SNIC Settings**: Select the interface for which you want to view the interface usage information. You can select one of the following from the drop-down list: **HA**, **LAN1**, or **LAN2**. You can view the interface usage information for the selected interface in the **Interface Usage** tab. This is displayed only when you select a Grid member that supports Infoblox Advanced DNS Protection.  
Note that you can select the HA port even though the Grid member is not an HA pair, because the HA port on a single member can be exposed to potential attacks.
- **NIC Settings**: Select the interface for which you want to view the interface usage information. You can select one of the following from the drop-down list: **HA**, **LAN1**, or **LAN2**. You can view the interface usage information for the

selected interface in the Interface Usage (LAN1) tab. This is displayed only when you select a Grid member that supports Software ADP.

- **Events Over Time:** Select the severity level, **Critical**, **Major**, **Warning**, or **Informational**, to view the details for a specific severity level. You can select one or all the available severity levels.
- Select the **Auto Refresh Period** checkbox to turn on auto-refresh, and specify the auto-refresh period in seconds. The default is 30 seconds.

#### **Warning**

If the *Detailed Status panel* is open, the following actions take place:

- Grid Manager auto refreshes at a rate of 30 seconds.
- Widgets that support user-specified auto refresh, refresh at the rate specified in the **Auto Refresh Period** field.

Therefore, even if you set the session timeout to be to greater than the auto refresh rate, auto refresh still takes place. The Grid Manager session does not time out because auto refresh takes precedence over the session timeout.

Click the Configure icon again to hide the configuration panel after you complete the modification. You can do the following in this widget:

- Click the **Summary** tab to view the statistics for the following resources in the format you selected:
- **Smart NIC CPU:** The percentage of Smart NIC CPU that is in use. This is displayed only when you select a Grid member that supports Infoblox Advanced DNS Protection.
- **Traffic being dropped:** The percentage of traffic dropped. It is displayed for both LAN1 and LAN2 interfaces.
- **Traffic being received:** The percentage of traffic received. It is displayed for both LAN1 and LAN2 interfaces.
- Click the **Events Over Time** tab to view information about the threat protection event counts for each severity level over the given time frame. It displays line graphs that show the threat protection event counts for each event severity over the last 30 minutes. Each event severity is represented by a different color line graph. You can hover your mouse over the graph to view the coordinates of any point in the graph. You can also click the Events Over Time legend and use it as a filter to view the graph for specific severity level.
- Click the **Top 10 Rules** tab to view information about the threat protection rules that have the most number of hits. It displays a bar chart to track the top 10 threat protection rules with the most number of hits for critical, major, and warning severity levels. Each event severity is displayed in a different color. If you have configured a Reporting member in the Grid, the **Go To History** link is displayed on this tab. You can click **Go To History** to view the *Threat Protection Top Rules Logged* report on the **Reporting** tab. To navigate back to the Security dashboard from the **Reporting** tab, click **Back to Security Dashboard** at the top left corner of the navigation bar in the **Reporting** tab.
- Click the **Top 10 Clients** tab to view information about the top sources (client IP addresses) that triggered threat protection rules. It displays a bar chart to track the top 10 clients with the most number of hits. If you have configured a Reporting member in the Grid, the **Go To History** link is displayed on this tab. You can click **Go To History** to view the *Threat Protection Top Rules Logged by Source* report on the **Reporting** tab. To navigate back to the Security dashboard from the **Reporting** tab, click **Back to Security Dashboard** at the top left corner of the navigation bar on the **Reporting** tab.
- Click the **Interface Usage** tab to view information about the interface usage (in megabytes per second) over a given time frame. It displays line graphs that show the interface usage trends for the selected interface over the last 30 minutes. You can hover your mouse over the graph to view the coordinates of any point in the graph.
- Click the **Smart NIC CPU** tab to view the information about the percentage of CPU usage over a given time frame. It displays line graphs that show the CPU usage trends over the last 30 minutes. You can hover your mouse over the graph to view the coordinates of any point in the graph. This is displayed only when you select a Grid member that supports Infoblox Advanced DNS Protection.

## Response Policy Zone (RPZ) Status for Grid

The *Response Policy Zone (RPZ) Status for Grid* widget provides statistical information about RPZ hits for the Grid. This widget contains the following tabs: **Top 10 Grid Members**, **RPZ Recent Hits**, **Trend** and **Health**.

You can do the following in this widget:



- Select a graph configuration, **Client Hits**, **Passthru Hits**, **Blocked Hits**, or **Substituted Hits**, to view details of a specific RPZ rule. You can select either one or all the available graph configurations. Note that **Client Hits** is displayed only when the graph type is **Line Diagram**.
- Select a graph type, **Stacked Diagram** or **Line Diagram**, to display data in the required diagrammatic format. This option is enabled only when you click the **Trend** tab and disabled when you click the **Top 10 Grid Members**, **RPZ Recent Hits**, or **Health** tabs. For more information, see [Trend](#) below.
- Click the **Top 10 Grid Members** tab to view information about the top 10 Grid members that have the most number of RPZ hits. For more information, see [Top 10 Grid Members](#) below.
- Click the **RPZ Recent Hits** tab to view information about the latest five RPZ hits with unique client addresses. For more information, see [RPZ Recent Hits](#) below.
- Click the **Trend** tab to view RPZ hit statistics for the Grid. For more information, see [Trend](#) below.
- Click the **Health** tab to view information about RPZ zones and their last updated times. For more information, see [Health](#) below.

Note that you must install the RPZ license and enable **RPZ logging** to access this widget. For more information about installing licenses and enabling RPZ logging, see [License Requirements and Admin Permissions](#) and [Using a Syslog Server](#).

### Top 10 Grid Members

The **Top 10 Grid Members** tab displays a stacked bar chart that tracks the top Grid members with the most total counts of RPZ hits. Each RPZ hit type is represented with a different color. The report displays the top 10 members in descending order.

### RPZ Recent Hits

The **RPZ Recent Hits** tab displays the data that is collected from the most recent hits of five unique clients, identified by their IP addresses, during the last 24 hours. NIOS retrieves this data from the syslog. This tab does not display any data when there are no syslog messages or if RPZ logging is disabled. NIOS displays an error message if RPZ logging is disabled. For more information about enabling RPZ logging and [Setting DNS Logging Categories](#), see [Using a Syslog Server](#).

Grid Manager retrieves recent hits from the Grid members. If a member has an RPZ license installed, then NIOS will parse the syslog every 60 seconds to collect the data. NIOS parses the generated data to identify the five most recent hits. It searches for these fields in the syslog message: CEF: data string(RPZ syslog) and src fields.

The NIOS appliance remembers the start and end time of previously searched operations to optimize the recent hits data collection, so that the same data is not searched again. Note that when the same client makes repeated queries in the last 24 hours, then there might be less than five unique client hits. You cannot sort or filter values on this tab.

This tab displays the following information:

- **Client IP Address:** IP address of the client that made the recent hits.
- **Requested FQDN:** The domain name or IP address that triggered the RPZ rule. For example, consider an RPZ rule `test.com.rpz.com`, which queries for `test.com`. In this example, `test.com` is the requested FQDN.
- **RPZ Entry:** The RPZ rule that queried a domain name or an IP address. In the above example, [test.com.rpz.com](#) is the RPZ rule.
- **Timestamp:** The date and time when the hit occurred.

Consider an example in which you query an RPZ zone and the NIOS appliance logs the following message in the syslog:

```
CEF:0|Infoblox|NIOS|6.9.0-219291|RPZ-QNAME|NODATA|4|app=DNS dst=10.35.101.14
src=10.36.0.251 spt=44460 view=_default qtype=A msg="rpz QNAME NODATA
rewrite w18.vg \[A\] via w18.vg.fireeye.com"
```

This tab displays information in the corresponding fields as follows:

Fields	Description
Client IP Address	Data is retrieved from the src field. Example: 10.36.0.251
Requested FQDN	It is retrieved from the data between the rewrite and [A] via fields. Example: w18.vg.
RPZ Entry	It is retrieved from the data after the via in msg field. Example: w18.vg.fireeye.com
Timestamp	This is listed in the syslog.

You can export data displayed in this tab by clicking the *Export* icon. For more information about exporting displaying data, see [Importing and Exporting Data using CSV Import](#).

## Trend

The **Trend** tab displays statistics of RPZ hits during the last 60 minutes for the Grid. You can use a stacked graph or a line graph to view the hits. Each of the RPZ policy is represented with a different color. This tab displays the following information:

- **Client Hits:** Total number of queries that triggered an RPZ policy. Note that this option is not displayed when you choose **Stacked Diagram**, but displayed only when you choose **Line Diagram**.
- **Passthru Hits:** Total number of queries that triggered a **Passthru RPZ** rule. For more information about passthru rules, see [Configuring Rules for RPZs](#).
- **Blocked Hits:** Total number of queries that triggered a **Block (No Data)** or **Block (No Such Domain)** RPZ rule. For more information about Managing Block (No Data) Rules or Managing Block (No Such Domain) Rules, see [Configuring Rules for RPZs](#).
- **Substitute Hits:** Total number of queries that triggered a **Substitute (Domain Name)** or **Substitute (Record)** RPZ rule. For more information, see [Managing Substitute \(Domain Name\) Rules](#) and [Managing Substitute \(Record\) Rules](#).
- **Timestamp:** The graph displays a 24 hours time window. Note the following about this tab:
  - The statistical data in DNS service will be reset when you stop and restart the DNS service or if you force an active DNS service to restart regardless of its state. This results in loss of prior data.
  - Using this graph, you can view the timestamp of statistics collection.

## Health

The **Health** tab displays information of RPZ zones and their last updated date and time. This data is retrieved directly from the database. Note that you cannot sort or filter values on this tab. You can export the data displayed on this tab by clicking the *Export* icon. For more information, see [Exporting Displayed Data](#).

## Response Policy Zone (RPZ) Status for Member

The *Response Policy Zone (RPZ) Status for Member* widget provides statistical information about RPZ hits for the selected member. This widget contains the following tabs: **RPZ Recent Hits**, **Trend**, and **Health**.

You can do the following in this widget:

- Click **Select Member**. In the *Member Selector* dialog box, choose a Grid member to view the RPZ hits, or statistics, or RPZ zones and their last updated date and time.
- Select a graph configuration, **Client Hits**, **Passthru Hits**, **Blocked Hits**, or **Substituted Hits**, to view details of a specific RPZ rule. You can select either one or all the available graph configurations. Note that **Client Hits** is displayed only when the graph type is **Line Diagram**.

- Select a graph type, **Stacked Diagram** or **Line Diagram**, to display data in the required diagrammatic format. This option is enabled only when you click the **Trend** tab and disabled when you click the **Top 10 Grid Members**, **RPZ Recent Hits**, or **Health** tabs. For more information, see [Trend](#) below.
- Click **View Syslog** to view the last 20 RPZ events that are logged in the syslog. For more information, see [Previewing the Syslog](#) below.
- Click the **RPZ Recent Hits** tab to view information about the latest five RPZ hits with unique client addresses. For more information, see [RPZ Recent Hits](#) below.
- Click the **Trend** tab to view RPZ hit statistics on the selected member. For more information, see [Trend](#) below.
- Click the **Health** tab to view information about RPZ zones and their last updated times. For more information, see [Health](#) below.

Note that you must install the RPZ license and enable **RPZ logging** to access this widget. For more information about installing licenses and enabling RPZ logging, see [License Requirements and Admin Permissions](#) and [Setting DNS Logging Categories](#).

### RPZ Recent Hits

The **RPZ Recent Hits** tab displays the data that is collected from the most recent hits of five unique clients, identified by their IP addresses, during the last 24 hours. NIOS retrieves this data from the syslog. This tab does not display any data when there are no syslog messages or if RPZ logging is disabled. NIOS displays an error message if RPZ logging is disabled. For more information about enabling RPZ logging, see [Setting DNS Logging Categories](#).

Grid Manager retrieves recent hits from the selected member. If a member has an RPZ license installed, then NIOS will parse the syslog every 60 seconds to collect the data. NIOS parses the generated data to identify the five most recent hits. It searches for these fields in the syslog message: CEF: data string(RPZ syslog) and src fields.

The NIOS appliance remembers the start and end time of previously searched operations to optimize the recent hits data collection, so that the same data is not searched again. Note that when the same client makes repeated queries in the last 24 hours, then there might be less than five unique client hits. You cannot sort or filter values on this tab.

This tab displays the following information:

- **Client IP Address:** IP address of the client that made the recent hits.
- **Requested FQDN:** The domain name or IP address that triggered the RPZ rule. For example, consider an RPZ rule [test.com.rpz.com](#), which queries for [test.com](#). In this example, [test.com](#) is the requested FQDN.
- **RPZ Entry:** The RPZ rule that queried a domain name or an IP address. In the above example, [test.com.rpz.com](#) is the RPZ rule.
- **Timestamp:** The date and time when the hit occurred.

Consider an example in which you query an RPZ zone and the NIOS appliance logs the following message in the syslog:

```
CEF:0|Infoblox|NIOS|6.9.0-219291|RPZ-QNAME|NODATA|4|app=DNS dst=10.35.101.14
src=10.36.0.251 spt=44460 view=_default qtype=A msg="rpz QNAME NODATA
rewrite w18.vg \[A\] via w18.vg.fireeye.com"
```

This tab displays information in the corresponding fields as follows:

Fields	Description
Client IP Address	Data is retrieved from the src field. Example: <a href="#">10.36.0.251</a>
Requested FQDN	It is retrieved from the data between the <code>rewrite</code> and <code>\[A\] via</code> fields. Example: <a href="#">w18.vg</a> .

Fields	Description
RPZ Entry	It is retrieved from the data after the <code>via</code> in <code>msg</code> field. Example: <a href="#">w18.vg.fireeye.com</a>
Timestamp	This is listed in the syslog.

## Trend

The **Trend** tab displays statistics of RPZ hits on the selected member during the last 60 minutes. You can use a stacked graph or a line graph to view the hits. DNS service generates RPZ statistics for the selected member. Each of the RPZ policy is represented with a different color. This tab displays the following information:

- **Client Hits:** Total number of queries that triggered an RPZ policy. Note that this option is not displayed when you choose **Stacked Diagram**, but displayed only when you choose **Line Diagram**.
- **Passthru Hits:** Total number of queries that triggered a **Passthru** RPZ rule. For more information about passthru rules, see [Managing Passthru Rules](#).
- **Blocked Hits:** Total number of queries that triggered a **Block (No Data)** or **Block (No Such Domain)** RPZ rule. For more information, see [Managing Block \(No Data\) Rules](#) or [Managing Block \(No Such Domain\) Rules](#) respectively.
- **Substituted Hits:** Total number of queries that triggered a **Substitute (Domain Name)** or **Substitute (Record)** RPZ rule. For more information, see [Managing Substitute \(Domain Name\) Rules](#) and [Managing Substitute \(Record\) Rules](#).
- **Timestamp:** The graph displays a 24 hours time window. Note the following about this tab:
  - The statistical data in DNS service will be reset when you stop and restart the DNS service or if you force an active DNS service to restart regardless of its state. This results in loss of prior data.
  - Using this graph, you can view the timestamp of statistics collection.

## Health

The **Health** tab displays information of RPZ zones on the selected member and their last updated date and time. This data is retrieved directly from the database. Note that you cannot sort or filter values on this tab. You can export the data displayed on this tab by clicking the *Export* icon.

## Previewing the Syslog

You can view the RPZ events that are logged in the syslog for a selected Grid member. Note that the preview displays only the last 20 RPZ events from the syslog. This wizard displays the following information:

- **Timestamp:** The date and time when the hit occurred.
- **Facility:** The location on the syslog server sorting the log message.
- **Level:** The severity of the message. This can be **ALERT**, **CRITICAL**, **DEBUG**, **EMERGENCY**, **ERROR**, **INFO**, **NOTICE**, or **WARNING**.
- **Server:** The name of the server that logged this message, plus the process ID.
- **Message:** Detailed information about the RPZ query. You can click the **Go to Syslog Viewer** link to view the RPZ events that are logged in the syslog. NIOS displays all the RPZ events that are logged in the syslog for the selected member and the **Quick Filter** is set to **RPZ Incident Logs** by default. For more information, see

## Tenant & VMs

The *Tenant & VMs* tab displays a table that shows the total number of **Tenants**, **Cloud VMs**, and **IP Addresses**, depending on your configuration. It also displays the average number of cloud VMs and IP addresses per tenant.

## Fixed vs. Floating

The *Fixed vs. Floating* tab displays IP address allocation for cloud objects. It displays the total number of fixed address allocation and floating address allocation, depending on your configuration. It also displays a pie chart indicating the percentage for each allocation.

## Available vs. Allocated

The *Available vs. Allocated* tab displays IP address allocation for available versus allocated IP addresses. It displays the total number of available IP addresses versus allocated IP addresses, depending on your configuration. It also displays a pie chart indicating the percentage for each allocation.

## Threat Analytics Status for Grid

The *Threat Analytics Status for Grid* widget displays the statistical information about the DNS tunneling events. This widget contains the following tabs: *Detections Over Time*, *Top 10 Grid Members*, and *Detections*. You can do the following in this widget:

### Warning

If the *Detailed Status panel* is open, the following actions take place:

- Grid Manager auto refreshes at a rate of 30 seconds.
- Widgets that support user-specified auto refresh, refresh at the rate specified in the **Auto Refresh Period** field.

Therefore, even if you set the session timeout to be to greater than the auto refresh rate, auto refresh still takes place. The Grid Manager session does not time out because auto refresh takes precedence over the session timeout. For more information about widgets, see [Status Dashboard](#).

- Turn on auto-refresh.  
Click the Configure icon, select the **AutoRefreshPeriod** checkbox, and specify the refresh period in seconds. The default auto refresh period is 30 seconds. Click the Configure icon again to hide the configuration panel after you complete the modification.
- Click the *Detections Over Time* tab to view information about the detected DNS tunneling events in a given time frame.
- Click the *Top 10 Grid Members* tab to view information about the top 10 Grid members with the most total counts of detections by type.
- Click the *Detections* tab to view information about all the detected DNS tunneling events.

## Detections Over Time

The **Detections Over Time** tab displays a line graph that tracks the number of detected DNS tunneling events over the given time frame. You can hover your mouse over the graph to view the coordinates of any point in the graph.

## Top 10 Grid Members

The **Top 10 Grid Members** tab displays a stacked bar chart that tracks the top Grid members with the most total counts of detected DNS tunneling events by type. The report displays the top 10 Grid members in descending order.

## Detections

The **Detections** tab displays information about all the detected DNS tunneling events. This tab displays the following information about each detection in table format:

- **Client IP Address:** The IP address of the client.
- **Domain:** The domain name of the client.
- **Timestamp:** The timestamp when the event occurred.

- **Module:** Displays the threat analytics module.

## Threat Analytics Status for Member

The *Threat Analytics Status for Member* widget displays statistics about the DNS tunneling events for a specific Grid member.

To configure the *Threat Analytics Status for Member* widget, click the Configure icon and complete the following:

### Warning

If the *Detailed Status panel* is open, the following actions take place:

- Grid Manager auto refreshes at a rate of 30 seconds.
- Widgets that support user-specified auto refresh, refresh at the rate specified in the **Auto Refresh Period** field.

Therefore, even if you set the session timeout to be to greater than the auto refresh rate, auto refresh still takes place. The Grid Manager session does not time out because auto refresh takes precedence over the session timeout. For more information about widgets, see [Status Dashboard](#).

- Click **Select Member** to select a Grid member. If there are multiple members, the *MemberSelector* dialog box is displayed, from which you can select a member. Click the required member name in the dialog box. You can also click **Clear** to clear the displayed member and select a new one.
- Select the **Auto Refresh Period** checkbox to turn on auto-refresh, and specify the auto-refresh period in seconds. The default is 30 seconds.

Click the Configure icon again to hide the configuration panel after you complete the modification.

You can do the following in this widget:

- Click the **Detections Over Time** tab to view information about the DNS tunneling event count for the selected Grid member in a given time frame. It displays a line graph that tracks the number of DNS tunneling event detections in a given time frame. You can hover your mouse over the graph to view the coordinates of any point in the graph.
- Click the **Detections** tab to view information about all the detected DNS tunneling events. This tab displays the following information in table format:
  - **Client IP Address:** The IP address of the client.
  - **Domain:** The domain name of the client.
  - **Timestamp:** The timestamp when the event occurred.
  - **Module:** Displays the threat analytics module. This tab displays only the last 15 detections.

## Pool Licenses Statistics

The *Pool Licenses Statistics* widget displays information about pool license allocation in your Grid. Pool licenses are dynamic service and feature licenses, such as vNIOs, DNS, DHCP, and Cloud Platform that you purchase for vNIOs and cloud deployments based on your evolving business needs. The Grid Master keeps track of dynamic licenses that are allocated to vNIOs members and adjusts the total number of available dynamic licenses for each feature and service. In the widget, you can select the license type and Grid Manager displays the total numbers of assigned and available licenses based on the selected license type.

To configure the *Pool Licenses Statistics* widget, click the Configure icon and complete the following:

- **Show Usage for:** From the drop-down list, select a license type for which you want the appliance to display dynamic license usage.

The widget displays license usage and the numbers of assigned and available licenses.

You can also view the total number of dynamic licenses installed for each feature and service, the number of active and available licenses, their usage, and other related information on the Grid tab -> Licenses tab -> Pool tab of Grid Manager.

For information about how to view dynamic licenses, see [Managing Licenses](#).

## DNS Record Scavenging

The DNS Record Scavenging dashboard widget displays statistics about the scavenging of stale DNS records. The widget displays the following information for the current or last known scavenging activities:

- **Status:** The status of the scavenging operation.
- **Start:** The start time of the scavenging operation.
- **End:** The end time of the scavenging operation.
- **User:** The user who initiated the scavenging operation.
- **Selected Object:** The Grid, view, or zone affected by record scavenging.
- **Action:** The action applied to the scavenging operation.
- **Processed Records:** The number of DNS records processed.
- **Reclaimable Records:** The number of DNS records marked as reclaimable.
- **Reclaimed Records:** The number of DNS records removed during the scavenging operation.

Click an icon on the filter panel, as illustrated in the Figure 2.2, to add a widget to the desired dashboard. Filter panel is categorized in to the following: Cloud , Security , DNS/DHCP , and Reset . When you click on an icon, Grid Manager displays thumbnails of the widgets belonging to the respective filter. If you click filters one after the other without clicking Reset, Grid Manager displays thumbnails of all widgets along with the icon that indicates the category to which the widget belongs. Click Reset to view only those widgets that belong to the selected category.

## IPAM Task Pack

The IPAM task pack contains the following tasks:

- Add Networks
- Add Hosts
- Add Fixed Addresses
- Add A Record
- Add CNAME Record
- Add TXT Record
- Add MX Record
- CSV Import

Depending on your administrative permissions and the Dashboard template configuration, Grid Manager displays tasks you can access in specific task packs. You can configure your task packs by adding or hiding certain tasks. For information about Dashboard templates, see [Configuring Dashboard Templates](#).

To hide tasks in a task pack:

1. Click the Configure icon at the upper right corner of the task pack.
2. In the configuration panel, select the tasks you want to hide from the Active Tasks table. You can use SHIFT+click and CTRL+click to select multiple tasks.
3. Click the left arrow to move the selected tasks to the Available Tasks table.

Click the Configuration icon again to hide the configuration panel after you complete the modification.

## Required Licenses for IPAM Tasks

The below Required Licenses for IPAM Tasks table lists the required licenses for viewing and performing IPAM tasks on the Tasks Dashboard.

### *Required Licenses for IPAM Tasks*

Task	Required Licenses
Add Networks	DHCP or MSMGMT license
Add Hosts	DNS or DHCP license



Task	Required Licenses
Add Fixed Addresses	DHCP or MSMGMT license
Add A Record	DNS or MSMGMT license
Add CNAME Record	DNS or MSMGMT license
Add TXT Record	DNS or MSMGMT license
Add MX Record	DNS or MSMGMT license

For information about how to install licenses, see [Managing Licenses](#).

## Add Networks

You can create IPv4 and IPv6 networks from the Tasks Dashboard (either from scratch or from a network template that contains predefined properties). You can also create networks from the **Data Management** tab. For more information about IPv4 and IPv6 networks, see [Configuring IPv4 Networks](#) and [Configuring IPv6 Networks](#).

To add networks from the Tasks Dashboard:

- Click **Add Networks** in the IPAM task pack and complete the following in the *Add Networks* wizard:
  - Regional Internet Registry:** This section appears only when support for RIR updates is enabled. For information about RIR, see [RIR Registration Updates](#). Complete the following to create an RIR IPv4 network container or network:
    - Internet Registry:** Select the RIR from the drop-down list. The default is **RIPE**. When you select **None**, the network is not associated with an RIR organization.
    - Organization ID:** Click **Select Organization** and select an organization from the *RIR Organization Selector* dialog box.
    - Registration Status:** The default is **Not Registered**. When adding an RIR allocated network, you can change this to **Registered** and select the **Do not update registrations** checkbox below. Note that when you select **API** as the communication method, the registration status will be updated automatically after the registration update is completed. However, when you select **Email** as the communication method, the registration status will not be automatically updated. If you are creating a new network and the registration update is completed successfully, the status will be changed to **Registered**. If the update fails, the status will be changed to **Not Registered**. The updated status and timestamp are displayed in the **Status of last update** field in the *IPv4/IPv6 Network Container* or *IPv4/IPv6 Network* editor.
    - Registration Action:** Select the registration action from the drop-down list. When you select **Create**, the appliance creates the IPv4 or IPv6 network and assigns it to the selected organization. When you select **None**, the appliance does not send registration updates to RIPE. When you are adding an existing RIR allocated network to NIOS, select **None**. When you are adding networks to an RIR allocated network (a parent network), select **Create**. Ensure that the parent network associated with an RIR organization already exists.
    - Do not update registrations:** Select this checkbox if you do not want the appliance to submit RIR updates to RIPE. By default, the appliance sends updates to the RIR database based on the configured communication method.
  - Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the network.
  - Protocol:** Select **IPv4** to add IPv4 networks and **IPv6** to add IPv6 networks.
  - Netmask:** Enter the netmask or use the netmask slider to select the appropriate number of subnet mask bits for the network.
  - Template:** Click **Select Template** to select a network template. When you use a template to create a network, the configuration of the template applies to the new network. If the template specifies a fixed



netmask, you cannot edit the netmask in this dialog. You can click **Clear** to remove the template. For information about templates, see [About IPv6 Network Templates](#).

- **Use Active Directory Sites:** This checkbox is displayed only if you install the Microsoft license. Click the Add icon to associate multiple Active Directory Sites with the network. When you click Add, the appliance displays the following:
  - **Active Directory Domain:** The Active Directory Domains that are synchronized from the Microsoft server.  
Click an Active Directory Domain that you want to associate.  
To search for a particular Active Directory Domain, specify the respective name and click **Go**. If there are multiple Active Directory Domains, the appliance displays the list of such domains by paging to the next page. You can use the page navigation buttons that are displayed at the bottom of this column to navigate through the Active Directory Domains. You can also refresh the values in the column using the Refresh icon.
  - **Active Directory Site:** The Active Directory Sites that are associated with the selected Active Directory Domain. Click an Active Directory Site that you want to associate with the network.  
To search for a particular Active Directory Site, specify the respective name and click **Go**. If there are multiple Active Directory Sites, the appliance displays the list of such sites by paging to the next page. You can use the page navigation buttons that are displayed at the bottom of this column to navigate through the Active Directory Sites. You can also refresh the values in the column using the Refresh icon.
  - Click **Add** to add the selected Active Directory Sites to the network or click **Cancel** to cancel the operation. The appliance displays these domains and sites in the respective columns. Click the x icon if you want to close the Active Directory Domains and Sites selector.
  - Click the Delete icon to delete Active Directory Sites that are associated with the network.

For more information about Active Directory Domains and Sites, see [Configuring Active Directory Sites and Services](#).

NIOS may execute discovery on the newly created network after you save your settings. When you create a network in NIOS, it inherits its discovery capabilities (whether or not it is immediately discovered, its polling settings, and any possible exclusions from discovery), from its parent network (if it has one) or its network container. If the new network is a parent network, it inherits its polling settings from the Grid and its discovery member selection and Enable Discovery action must be defined by the user.

- **Networks:** Do one of the following to add new networks:

Click the Add icon to create a new network.

- For IPv4 networks: Grid Manager adds a row to the table. Enter the network address in the **Network** field. Click the Add icon to add another network. You can also select a network and click the Delete icon to delete it.
- For IPv6 networks: If you are adding a network for a previously defined global IPv6 prefix, you can select the prefix from the **IPv6 Prefix** drop-down list. The default is **None**, which means that you are not creating an IPv6 network for a previously defined subnet route. If you have defined a global prefix at the Grid level, the default is the global prefix value. Click **Add** and Grid Manager adds a row to the table. Enter the network address in the **Network** field. When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2001:0db8:0000:0000:0000:0000:0102:0304 can be shortened to 2001:db8::0102:0304. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered. Click **Add** again to add another network. You can also select a network and click the Delete icon to delete it.

or

Click the Next Available icon to have the appliance search for the next available network. For more information about the next available network, see [Configuring the Next Available Network or IP Address](#). Complete the following in the Next Available Networks section:

- **Create new network(s) under:** Enter the network container in which you want to create the new network. When you enter a network that does not exist, the appliance adds it as a network container. When you enter a network that is part of a parent network, the parent network is converted into a network container if it does not have a member assignment or does not contain address ranges, fixed addresses, reservations, shared networks, and host records that are served by DHCP. When you enter a network that has a lower CIDR than an existing network, the appliance creates the network as a parent network and

displays a message indicating that the newly created network overlaps an existing network. You can also click **Select Network** to select a specific network in the *Network Selector* dialog box. For information about how the appliance searches for the next available network, see [Configuring the Next Available Network or IP Address](#).

- **Number of new networks:** Enter the number of networks you want to add to the selected network container. Note that if there is not enough network space in the selected network to create the number of networks specified here, Grid Manager displays an error message. The maximum number is 20 at a time. Note that when you have existing networks in the table and you select one, the number you enter here includes the selected network.
  - Click **Add Next** to add the networks. Grid Manager lists the networks in the table. You can click **Cancel** to reset the values.

Note that you must click Add Next to add the network container you enter in the Next Available Networks section. If you enter a network in the Next Available Networks section and then use the Add icon to add another network, the appliance does not save the network you enter in the Next Available Networks section until you click Add Next.

- **Extensible Attributes:** Click the Add icon to enter extensible attributes. Grid Manager adds a row to the table each time you click the Add icon. Select the row and the attribute name from the drop-down list, and then enter the value. All inheritance attributes which can be inherited from a parent object will be automatically inherited when you add a network. Inheritable extensible attributes that are required are automatically displayed. Optional extensible attributes that are not inheritable are not automatically displayed. For more information about extensible attributes, see [Managing Extensible Attributes](#).
- If you are adding an RIR network, the RIR network attribute table appears. For information about these attributes and how to enter them, see [Managing RIR Attributes](#).

**Preview RIR Submissions:** Click this to view the updates before the appliance submits them to the RIPE database. This button is enabled only when the registration action is **Create**, **Modify**, or **Delete**, and the **Do not update registrations** checkbox is not selected.

2. Save the configuration.  
or

Click the **Schedule** icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information, see [Managing Extensible Attributes](#).

The appliance saves the networks you just created, and Grid Manager displays them in the Result page. When you click a newly created network on this page, Grid Manager displays the **IP Map** panel from which you can view detailed information about the network. For information about the IP Map panel, see [Viewing and Managing IPv4 Addresses](#).

You can also add and modify other information about the networks you just created. For information about modifying network information, see [Managing IPv4 DHCP Data](#) and [Managing IPv6 DHCP Data](#).

## Add Hosts

Host records provide a unique approach to the management of DNS, DHCP, and IPAM data. By using host records, you can manage multiple DNS records and DHCP and IPAM data collectively, as one object on the appliance. You can add IPv4 and IPv6 addresses to host records from the Tasks Dashboard or the **Data Management** tab. Note that when you add a host record from the Tasks Dashboard, they are configured only for DNS. For more information about Infoblox host records, see [About Host Records](#).

To add host records from the Tasks Dashboard:

1. Click **Add Hosts** in the IPAM Task Pack and complete the following in the *Add Hosts* wizard:
  - **Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the host record.
  - **Zone Name:** Click **Select** to select a DNS zone from the *Zone Selector* dialog box.
  - **Exclude from Network Discovery and Immediate Discovery:** When creating the new Host record, you can direct NIOS to immediately discover the host, or to exclude it from network discovery. By default, the Add Hosts task enables immediate discovery.
  - **DNS View:** Displays the DNS view of the selected zone.
  - **Hosts:** Do one of the following to add a host record:

Click the Add icon and the appliance adds a row to the table. Complete the following in the table to add a new host record:

- **Name:** Enter the name of the host record.
- **Zone:** Displays the DNS zone you select in **Zone Name**. When you enter a different zone here, the appliance displays an error message.
- **Address:** Enter the IP address you want to associate with this host record.

or

Click the Next Available icon to have the appliance search for the next available IP address for the host record. For information about the next available IP address, see [Configuring the Next Available Network or IP Address](#). Complete the following in the Next Available IP section:

- **Create new host addresses under:** Click **Select** to select the network or address range in the *Network/Range Selector* dialog box from which you want the appliance to search for the next available IP address for this host record.
- **Number of new host addresses:** Enter the number of host addresses. Note that if there is not enough space in the selected network or address range to create the number of host addresses specified here, Grid Manager displays an error message. The maximum number is 20 at a time. Note that when you have existing host addresses in the table and you select one, the number you enter here includes the selected host address.
- Click **Add Next** to add the IP addresses to their corresponding hosts. Grid Manager lists the host addresses in the table. Ensure that you enter a name for each host record.
- **Extensible Attributes**
  - **Apply to all above hosts:** Select this to associate extensible attributes with all hosts that you have defined. This is selected by default. You can define and associate multiple extensible attributes with multiple hosts at once.
  - **Apply to selected host:** Select this to associate extensible attributes with the selected host only. Note that when you select this option for another host in the list, the **Extensible Attributes table** is refreshed for you to associate a different set of extensible attributes with the selected host.
  - **Extensible Attributes table:** Click the Add icon to enter extensible attributes. The appliance adds a row to the table each time you click the Add icon. Select the row and the attribute name from the drop-down list, and then enter the value. All inheritance attributes which can be inherited from a parent object will be automatically inherited when you add a host. Inheritable extensible attributes that are required are automatically displayed. Optional extensible attributes that are not inheritable are not automatically displayed. For more information about extensible attributes, see [Managing Extensible Attributes](#).

2. Save the configuration. or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information about how to schedule a task, see [Scheduling Tasks](#).

The appliance saves the host records you just created, and Grid Manager displays them in the Result page. When you click a newly created host on this page, Grid Manager displays the **Data Management -> DNS -> Zones** tab from which you can view information about the host record.

You can also add and modify other information about the host records. For information about modifying host information, see [About Host Records](#).

## Add Fixed Addresses

You can add IPv4 and IPv6 fixed addresses from the Tasks Dashboard or from the **Data Management** tab. For more information about fixed addresses, see [Configuring IPv4 Fixed Addresses](#) and [Configuring IPv6 Fixed Addresses](#).

To add fixed addresses from the Tasks Dashboard:

1. Click **Add Fixed Addresses** in the IPAM task pack and complete the following in the *Add Fixed Addresses* wizard:
  - **Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the fixed address.
  - **Protocol:** Select **IPv4** to add IPv4 addresses and **IPv6** to add IPv6 addresses.

- **Template:** Click **Select Template** to select a fixed address template. When you use a template to create a fixed address, the configuration of the template applies to the new fixed address. You can also click **Clear** to remove the template. For information about templates, see [Managing DHCP Templates](#).
- **Exclude from Network Discovery and Immediate Discovery.** When creating the new fixed address, you can direct NIOS to immediately discover the device associated with the fixed address, or to exclude it from network discovery. By default, the Add Fixed Addresses task enables immediate discovery.
- **Addresses:** Do one of the following to add fixed addresses:

Click the Add icon and Grid Manager adds a row to the table. Complete the following to create fixed addresses:

- For IPv4 fixed addresses: Enter the IPv4 address and MAC address. Click the Add icon to add another fixed address.
- For IPv6 fixed addresses: Enter the IPv6 address and DUID. Click the Add icon again to add another fixed address.

or

Click the Next Available icon to have the appliance search for the next available address. Complete the following:

- **Create new fixed addresses under:** Click **Select** to select the network or address range in the *Network/Range Selector* dialog box from which you want the appliance to search for the next available IP address for this fixed address.
- **Number of new fixed addresses:** Enter the number of fixed addresses you want to add to the selected network or address range. Note that if there is not enough space in the selected network or address range to create the number of fixed addresses specified here, Grid Manager displays an error message. The maximum number is 20 at a time. Note that when you have existing fixed addresses in the table and you select one, the number you enter here includes the selected fixed address.
- Click **Add Next** to add the fixed addresses. The appliance lists the fixed addresses to the table. Ensure that you enter the MAC address or DUID for each fixed address.
- **Extensible Attributes** table: Click the Add icon to enter extensible attributes. The appliance adds a row to the table each time you click the Add icon. Select the row and the attribute name from the drop-down list, and then enter the value. All inheritance attributes which can be inherited from a parent object will be automatically inherited when you add a fixed address. Inheritable extensible attributes that are required are automatically displayed. Optional extensible attributes that are not inheritable are not automatically displayed. For more information about extensible attributes, see [Managing Extensible Attributes](#).

2. Save the configuration.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

The appliance saves the fixed addresses you just created, and Grid Manager displays them in the Result page. When you click a newly created fixed address on this page, Grid Manager displays the **Data Management -> IPAM -> IP Map** or **List** tab from which you can view information about the fixed address.

You can also add and modify other information about the fixed addresses you just created. For more information about modifying fixed address information, see [Managing IPv4 DHCP Data](#) and [Managing IPv6 DHCP Data](#).

## Add A Record

An A (address) record is a DNS resource record that maps a domain name to an IPv4 address. You can add an A record from the Tasks Dashboard or from the **Data Management** tab. For more information about managing A records, see [Managing Resource Records](#).

To add networks from the Tasks Dashboard:

1. Click **Add A Record** in the IPAM task pack and complete the following in the *Add A Record* wizard:
2. In the *Add A Record* wizard, do the following:
  - **Name:** If Grid Manager displays a zone name, enter the hostname that you want to map to an IP address. The displayed zone name can either be the last selected zone or the zone from which you are adding the host record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name

in the dialog box and then enter the hostname. The name you enter is prefixed to the DNS zone name that is displayed, and the complete name becomes the FQDN (fully qualified domain name) of the host. For example, if the zone name displayed is corpxyz.com and you enter admin, then the FQDN becomes admin.corpxyz.com. Ensure that the domain name you enter complies with the hostname restriction policy defined for the zone. To create a wildcard A record, enter an asterisk \* in this field.

- **DNS View:** This field displays the DNS view to which the DNS zone belongs.
  - **Shared Record Group:** This field appears only when you are creating a shared record from the **Data Management** tab. Click **Select Shared Record Group**. If you have only one shared record group, the appliance displays the name of the shared record group here. If you have multiple shared record groups, select the shared record group in the *Shared Record Group Selector* dialog box. You can use filters or the **Go to** function to narrow down the list.
  - **Hostname Policy:** Displays the hostname policy of the zone.
  - In the **IP Addresses** section, click the Add icon and do one of the following:
    - Select **Add Address** to enter the IPv4 address to which you want the domain name to map.  
or
    - Select **Next Available IPv4** to retrieve the next available IP address in a network. If the A record is in zone that has associated networks, the *Network Selector* dialog box lists the associated networks. If the zone has no network associations, the *Network Selector* dialog box lists the available networks. When you select a network, Grid Manager retrieves the next available IP address in that network.
    - **Comment:** Optionally, enter additional information about the A record.
    - **Create associated PTR record:** Select this option to automatically generate a PTR record that maps the specified IP address to the hostname. To create the PTR record, the reverse-mapping zone must be in the database.
    - **Disable:** Select this checkbox to disable the record. Clear the checkbox to enable it.
3. Click **Next** to define extensible attributes. For information, see [Managing Extensible Attributes](#).
  4. Save the configuration or click the Schedule icon at the top of the wizard to schedule this task. For information about how to schedule a task, see [Scheduling Tasks](#).
  5. Click **Restart** if it appears at the top of the screen.

## Add CNAME Record

A CNAME record maps an alias to a canonical name. You can use CNAME records in both IPv4 forward- and IPv4 reverse-mapping zones to serve two different purposes. (At this time, you cannot use CNAME records with IPv6 reverse-mapping zones.) For more information about CNAME records, see [Managing CNAME Records](#).

To add a CNAME record from the Tasks Dashboard:

1. Click **Add CNAME Record** in the IPAM task pack and complete the following in the *Add CNAME Record* wizard:
  - **Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the CNAME record.
  - **Alias:** Click **Select Zone** to select a DNS zone from the *Zone Selector* dialog box. If you have only one zone, Grid Manager displays the zone name here when you click **Select Zone**. Enter the alias for the canonical name. For an IPv4 reverse-mapping zone, enter the host portion of an IP address. For example, if the full IP address is 10.1.1.1 in a network with a 25-bit netmask, enter 1. (The 10.1.1.0/25 network contains host addresses from 10.1.1.1 to 10.1.1.126. The network address is 10.1.1.0, and the broadcast address is 10.1.1.127.)
  - **DNS View:** Displays the DNS view of the selected zone.
  - **Shared Record Group:** This field appears only when you are creating a shared record from the **Data Management** tab. Click **Select Shared Record Group**. If you have only one shared record group, the appliance displays the name of the shared record group here. If you have multiple shared record groups, select the shared record group in the *Shared Record Group Selector* dialog box. You can use filters or the **Go to** function to narrow down the list.
  - **Canonical Name:** This field displays the domain name of either the current zone or the last selected zone. To add a CNAME record to a forward-mapping zone, enter the complete canonical (or official) name of the host. To add a CNAME record to a reverse-mapping zone, enter *host\_ip\_addr.prefix.network.in-addr.arpa* (host IP address + 2317 prefix + network IP address + in-addr.arpa). For example, enter 1.0.25.1.1.10.in-addr.arpa. This IP address must match the address defined in the PTR record in the delegated child zone.
  - **Comments:** Enter useful information about this record.



- **Disable:** Select the checkbox to disable the record without deleting its configuration. Clear the checkbox to enable the record.

2. Save the configuration, or click **Next** to define extensible attributes. For information about extensible attributes, see [Managing Extensible Attributes](#).

or

3. Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

Click **Restart** if it appears at the top of the screen.

## Add TXT Record

A TXT (text record) record contains supplemental information for a host. For example, if you have a sales server that serves only North America, you can create a text record stating this fact. You can create more than one text record for a domain name. You can add a TXT record from the Tasks Dashboard or the **Data Management** tab. For more information about TXT records, see [Managing TXT Records](#).

To add TXT records from the Tasks Dashboard:

1. Click **Add TXT Record** in the IPAM task pack and complete the following in the *Add TXT Record* wizard:

- **Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the TXT record.
- **Name:** If Grid Manager displays a zone name, enter the name to define a TXT record for a host or subdomain. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box. Then, enter the TXT record name. The appliance prefixes the name you enter to the domain name of the selected zone. For example, if you want to create a TXT record for a web server whose host name is `www2.corpxyz.com` and you define the TXT record in the `corpxyz.com` zone, enter **www2** in this field. To define a TXT record for a domain whose name matches the selected zone, leave this field empty. The appliance automatically adds the domain name (the same as the zone name) to the TXT record. For example, if you want to create a TXT record for the `corpxyz.com` domain and you selected the `corpxyz.com` zone, leave this field empty.
- **DNS View:** Displays the DNS view of the selected zone.
- **Shared Record Group:** This field appears only when you are creating a shared record from the **Data Management** tab. Click **Select Shared Record Group**. If you have only one shared record group, the appliance displays the name of the shared record group here. If you have multiple shared record groups, select the shared record group in the *Shared Record Group Selector* dialog box. You can use filters or the **Go to** function to narrow down the list.
- **Text:** Enter the text that you want to associate with the record. It can contain substrings of up to 255 bytes, up to a total of 512 bytes. Additionally, if you enter leading, trailing, or embedded spaces in the text, add quotes around the text to preserve the spaces. For example: `" v=spf1 include:corp200.com -all "`
- **Comments:** Enter useful information about this record.
- **Disable:** Select the checkbox to disable the record without deleting its configuration. Clear the checkbox to enable the record.

2. Save the configuration, or click **Next** to define extensible attributes. For information, see [Managing Extensible Attributes](#).

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

3. Click **Restart** if it appears at the top of the screen.



### Note

In a signed zone, if you add or delete a TXT record that has @ in its host name, you must also create or delete the corresponding NSEC3 record.

## Add MX Record

An MX (mail exchanger) record maps a domain name to a mail exchanger. A mail exchanger is a server that either delivers or forwards mail. You can specify one or more mail exchangers for a zone, as well as the preference for using each mail exchanger. A standard MX record applies to a particular domain or subdomain. You can add an MX record from the Tasks Dashboard or the **Data Management** tab. For more information about MX records, see [Managing MX Records](#).

To add MX records from the Tasks Dashboard:

1. Click **Add MX Record** in the IPAM task pack and complete the following in the *Add TXT Record* wizard:
  - **Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the MX record.
  - **Mail Destination:** If Grid Manager displays a zone name, enter the mail destination here. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *ZoneSelector* dialog box. Click a zone name in the dialog box, and then enter the mail destination. If you want to define an MX record for a domain whose name matches the zone you selected, leave this field blank. Grid Manager automatically adds the domain name (the same as the zone name) to the MX record. For example, if you want to create an MX record for a mail exchanger serving the corpxyz.com domain and you selected the corpxyz.com zone, and leave this field empty. If you want to define an MX record for a subdomain, enter the subdomain name. The appliance prefixes the name you enter to the domain name of the selected zone. For example, if you want to create an MX record for a mail exchanger serving site1.corpxyz.com—a subdomain of corpxyz.com—and you define the MX record in the corpxyz.com zone, enter site1 in this field. If you want to define an MX record for a domain and all its subdomains, enter an asterisk ( \* ) to create a wildcard MX record.
  - **DNS View:** Displays the DNS view of the selected zone.
  - **Shared Record Group:** This field appears only when you are creating a shared record from the **Data Management** tab. Click **Select Shared Record Group**. If you have only one shared record group, the appliance displays the name of the shared record group here. If you have multiple shared record groups, select the shared record group in the *Shared Record Group Selector* dialog box. You can use filters or the **Goto** function to narrow down the list.
  - **Host Name Policy:** Displays the hostname policy of the selected zone. Ensure that the hostname you enter complies with the hostname restriction policy defined for the zone.
  - **Mail Exchanger:** Enter the fully qualified domain name of the mail exchanger. You can even enter a dot character in this field (.). When you enter a dot, it means that the domain is a parked domain and will not receive or send email.
  - **Preference:** Select an integer from 10 to 100, or enter a value from 0 to 65535. The preference determines the order in which a client attempts to contact the target mail exchanger with 0 being the highest preference.
  - **Comment:** Enter useful information about this record.
  - **Disable:** Select the checkbox to disable the record without deleting its configuration. Clear the checkbox to enable the record.

2. Save the configuration, or click **Next** to define extensible attributes. For information, see [Managing Extensible Attributes](#).

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, click **Later** and enter a date, time, and time zone. For information about scheduling tasks, see [Scheduling Tasks](#).

3. Click **Restart** if it appears at the top of the screen.

## CSV Import

You can access **CSV Import Manager** and perform CSV imports, manage import jobs, and view import status. You can perform CSV imports from the Task Dashboard and the Toolbar. You can also click **CSV Import Manager**, which allows for importing of data, managing import jobs, and viewing import status.

You can click **New CSV import job** icon in the **CSV Job Manager** wizard to import a CSV file to the database. You can also verify the content in your CSV file before replacing the content of the database with the content in the imported CSV file. For detailed information about the CSV import feature, see [Importing and Exporting Data using CSV Import](#).

## NetMRI Task Pack

The NetMRI Tasks pack or, alternatively, Automation Tasks pack, requires the configuration, licensing and connection of an Infoblox NetMRI appliance to support automation tasks.

This topic describes the following:

- Enabling NetMRI Tasks
- Disabling NetMRI Tasks
- Registering NetMRI with NIOS
- Running NetMRI Tasks
- Viewing Tasks Execution Logs and Approving Tasks in the Task Viewer
- Viewing Tasks Execution Logs in NetMRI

### Enabling NetMRI Tasks

After you install the NetMRI appliance into a managed network and ensure the appliance is reachable by the NIOS Grid Master, register the NetMRI appliance with NIOS.

1. Click the Configure icon at the top right corner of the **Tasks** page.
2. Choose **Enable NetMRI Tasks**.  
The Enable NetMRI Tasks dialog box appears, requesting verification of your action:  
**Though you can see the change immediately, other users who are currently logged in will not see the change until they log in again.**  
**Are you sure you want to proceed?**
3. Click **Yes** to enable the NetMRI Tasks set.  
After a moment, the NetMRI Tasks panel appears.



#### Note

If two superusers are logged in to the NIOS system and one superuser enables the NetMRI Tasks pack on their console, the other superuser will not see the task pack on their console until their next login; the **Disable NetMRI Tasks** from the Configure icon menu shows the correct state.

### Disabling NetMRI Tasks

To disable the NetMRI Tasks pack:

1. Click the Configure icon at the top right corner of the Tasks page.
2. Choose **Disable NetMRI Tasks**.  
The Disable NetMRI Tasks dialog box appears, requesting verification of your action:  
**Though you can see the change immediately, other users who are currently logged in will not see the change until they log in again.**  
**Are you sure you want to proceed?**
3. Click **Yes** to disable the NetMRI Tasks pack.

### Registering NetMRI with NIOS

You must register a NetMRI appliance with NIOS using Grid Manager to support NetMRI Tasks. You need an account with admin privileges on the NetMRI appliance and the appliance hostname or IP address.





#### Note

Ensure that the user account that you use for the registration and further communication between the products is identical to the existing valid account on the NetMRI appliance.

To register NetMRI with NIOS:

1. From the **Dashboards** tab, select the **Tasks** tab.
2. At the top right corner of the **Automation Tasks** panel, click the down arrow icon -> **NetMRI Registration**.
3. In the **NetMRI Registration** dialog, do the following:
  - a. Enter the IP address or resolved host name of the NetMRI appliance supporting the Automation Tasks pack.  
Note that when you register NetMRI with a NIOS HA pair, you can register only one interface at a time. Use the IP address of the LAN1 interface, not the VIP address, for registration. When an HA failover occurs, the NetMRI registration is disabled. You can register the NetMRI appliance again after the failover.
  - b. Enter the **Admin Password**.
4. Click **Register**.

After registration, the **NetMRI Registration** menu item changes to read **NetMRI Deregistration** to support disconnecting from the NetMRI appliance.



#### Note

After you successfully register a NetMRI appliance with NIOS, you can use the Ecosystem > Cisco ISE Endpoint feature. It is available with the NetMRI license or Network Insight license that is installed by default on the discovery member in NIOS installations. This feature enables you to enhance identity management across devices and applications that are connected to your network routers and switches. You can monitor domain users, the IP addresses they log on to, the login status, and the time duration of their current status in the **IPAM** tab. For information about how to collect user and device information from Cisco ISE, see [Integrating Cisco ISE into NIOS](#).

## Running NetMRI Tasks

The NetMRI task pack contains the following tasks:

- Network Provisioning
- Port Activation
- VLAN Reassignment
- Bare Metal Provisioning
- Rogue DHCP Server Remediation

Depending on your administrative permissions, Grid Manager displays tasks that you can access in the **Automation tasks** panel. You can configure the displayed tasks pack by adding or hiding tasks.

To hide tasks:

1. Click the Configure icon at the upper right corner of the **Automation Tasks** panel.
2. In the **Active Tasks** pane, select the tasks you want to hide from the pack. You can use SHIFT+click and CTRL+click to select multiple tasks.
3. Click the left arrow to move the selected tasks to the **Available Tasks** pane.
4. Click the Configuration icon again to hide the panes.

## NetMRI Task Options

Tasks allow the assignment of job scripts to change and expand task functionality. These scripts reside on the NetMRI appliance and must be readable by the NIOS system to run the automation tasks. You can also select different scripts to

execute for automation tasks that provide that feature in NIOS. Three NetMRI tasks allow for the choosing of non-default scripts for task operation:

- Network Provisioning Task
- Port Activation Automation Task
- Specifying a Port Activation Script

## Network Provisioning

The Network Provisioning task runs in two modes: a basic mode with a much shorter list of configuration options, and a more complex mode that provides detailed configuration for provisioning a network, including the use of NIOS network views, extensible attributes and network templates.

New networks can be provisioned on routed networks and on switched networks. In the latter case, you can specify the new VLAN number and VLAN name for provisioning, along with the Device Group Device and Interface. The Device Group values are taken from the Device Groups defined on the NetMRI appliance from which NIOS obtains its data.

Network Provisioning supports two types of networks: IPv4, in which the new network is IPv4 only, and IPv4 and IPv6, in which the new network runs both protocol stacks.

## Simple vs. Complex Provisioning

Use of a Network View determines whether you use the simple or detailed views of provisioning a network. A network view is a single routing domain with its own networks and shared networks. In NIOS, all networks must belong to a network view. You can manage networks in one network view independently of other network views. Because network views are mutually exclusive, the networks in each view can have overlapping address spaces with multiple duplicate IP addresses without impacting network integrity.

Also, the same network segment can be present in multiple network views. When you create a new network, you select one view in which to place it, and preserve those values to apply to another view.

You also have the option to provision a single network segment without recourse to NIOS network views. The simple network provisioning option (accessible by simply clicking the IPv4 tool at the top of the Network Provisioning dialog box) allows you to specify as few as three values to configure a network.

The NIOS system also provides a **default** network view, which appears as an option for network provisioning.

If a single network view is configured in NIOS, you will not see a **Network View** option in the Network Provisioning task.

## Applying Extensible Attributes

Extensible attributes are associated with a specific network view, and are referenced by the Network Provisioning task. Should you configure a new network using a network view, you may need to consider the application of extensible attributes to the new network (they are not automatically applied, but will appear in the Network Provisioning dialog if those attributes are defined in the chosen Network View). Extensible attributes are generally defined for descriptive and tracking purposes in the network. A network view may have attributes such as Building, Country, Region, Site, State or VLAN, for example. Attributes are defined for network views in NIOS but are not defined by the NetMRI appliance.

If the NIOS system supports only a single network view, no View selections are made for the purposes of network provisioning.

To perform an automatic network provisioning task:

1. From the **Dashboards** tab, select the **Tasks** tab -> **Network Provisioning**.
2. Select the network type for provisioning: **IPv4** or **IPv4 and IPv6**.
3. To configure IPv4 provisioning:
  - a. Enter the required name value in the **Interface Hostname** field. (Examples include "eth0" and "serial0.")
  - b. Select the **DNS Zone** under which the hostname operates.
  - c. Choose a device group from the **Device Group** drop-down list.
  - d. From the Device drop-down list, choose the switch or router on which the network will originate.
  - e. If the selected device is a router, the **VLAN Number** and **VLAN Name** fields will be disabled.

- f. From the **Interface** list, choose the interface to which the network will be reassigned. The drop-down list contains all the interfaces from the chosen network device, and also shows the ports' respective states (up/down, up/up and so on).
- g. If the chosen device is a switch, enter the new **VLAN Number** on which the new network segment runs.
- h. If the chosen device is a switch, enter the new **VLAN Name** on which the new network segment runs.
- i. Click **Provision Network** to commit settings.
- j. Enter the **Parent Network** value (or click Select Network to choose the parent network from a list if using a Network View).
- k. Choose the **Network Template** from the drop-down list if one is provided by the chosen Network View. The Network template is otherwise optional.
  - l. The Provision Network task provides subnetting tools.
- m. Drag the **Netmask** slider to the required CIDR mask bit depth (1-32).
- n. In the **New Network** field, enter the IP prefix for the new network.
- o. In the **Router Address** field, enter the IP address for the router interface.
- p. Select any **Extensible Attributes** in the list if they are provided; otherwise, you can create new ones by clicking Add and choosing the **Attribute Name**, **Value** and the **Required** setting.

#### 4. To configure IPv6 provisioning:

- a. Enter the **Parent Network** value. Or, if using a Network View, click **Select Network** to choose the parent network from the list.
- b. Choose the **Network Template** from the drop-down list if one is provided by the chosen Network View. The Network template is otherwise optional.
 

The Provision Network task provides subnetting tools.
- c. **Drag the Netmask** slider to the required CIDR mask bit depth (1-32).
- d. In the **New Network** field, enter the IP prefix for the new network.
- e. In the **Router Address** field, enter the IP address for the router interface.
- f. Select any **Extensible Attributes** in the list if they are provided; otherwise, you can create new ones by clicking Add and choosing the **Attribute Name**, **Value** and the **Required** setting.

5. Enter the required name value in the **Interface Hostname** field. (Examples include "eth0" and "serial0.")

6. Select the **DNS Zone** under which the hostname operates.

7. Choose a device group from the **Device Group** drop-down list.

8. From the Device drop-down list, choose the switch or router on which the network will originate.

9. If the selected device is a router, the **VLAN Number** and **VLAN Name** fields will be disabled.

10. From the **Interface** list, choose the interface to which the network will be reassigned. The drop-down list contains all the interfaces from the chosen network device, and also shows the ports' respective states (up/down, up/up and so on).

If an interface shows **Routed** or **Switched**, it cannot be selected for provisioning as it is already being used as part of an active network.

11. If the chosen device is a switch, enter the new **VLAN Number** on which the new network segment runs.

12. If the chosen device is a switch, enter the new **VLAN Name** on which the new network segment runs.

13. Click **Provision Network** to commit settings.

The system sends the configuration request to the NetMRI appliance and displays the task configuration sequence.

#### Defining Options for the Network Provisioning Task

The Network Provisioning task provides several configuration options that affect how the task operates.

Hostname provisioning for interfaces is useful for troubleshooting purposes in the network, usually to ensure that an admin knows which router interface they are connecting through to communicate with the device. The hostname value is actually provisioned from within the Network Provisioning task. Enabling the Hostname Required? checkbox sets the NetMRI appliance to provision the network with hostnames applied to the router interfaces for easier identification.

Network provisioning requires that the system know exactly which IP address the gateway for the network will reside. For provisioning most networks, an Offset value of 1 indicates that the provisioned network gateway IP address ends with the host address of ...1, as in 192.168.1.1. **An Offset value of 1 will be by far the most common value for provisioning networks. Specifying an offset value other than 1 indicates that the gateway IP is a specified number of host values from the prefix address of the network. For example, setting an IPv4 Gateway Address Offset of 12 indicates that the IP for the**

gateway ends in \*.\*.12, as in 10.1.1.12. Offsets work the same way for any size network: for an example such as 10.1.1.64/26, and an offset of 12, the provisioned gateway IP would be 10.1.1.76.



#### Note

Make sure the defined offset value lies within the addressable boundaries of the provisioned network.

The same principles also apply for IPv6 networks, except that the IPv6 value is entered manually in hexadecimal instead of being selected from a drop-down list. Most provisioned IPv6 networks will use a /64 network address.

You can also select a different script from the default for the Network Provisioning task. To define settings for the Network Provisioning task:

1. From the **Dashboards** tab, select the **Tasks** tab. Under the **Network Provisioning** task, click the settings icon on the top right.
2. If the provisioning process requires a hostname, enable the **Hostname Required?** checkbox. (The network interface hostname ("eth0," "serial0") and the Zone that it belongs to are defined in the Network Provisioning task.)
3. Choose a gateway offset value from the **IPv4 Gateway Address Offset** drop-down list. If no value is selected, the offset value defaults to 1 for the provisioned network address.
4. If an IPv6 offset is required for provisioning an IPv6 network or for provisioning a network that supports both IPv4 and IPv6 addressing, enter the **IPv6 Gateway Address Offset** value in hexadecimal. If no value is entered, the offset value defaults to 0000.0000.0000.0001 for the provisioned network address, indicating an offset value of 1 for the gateway IP address.
5. In the **Script Name** dropdown, choose the script that you wish to run for the Port Activation task. The scripts are located on the Trinzi Automation 4000 appliance, and referenced for use by NIOS. By default, the bundled **Port Activation** script is selected.
6. Click **Save**.
7. Click **Cancel** to close the dialog.

The system sends the request to the NetMRI appliance and displays a **Provisioning Network Config updated** notification message.

## Port Activation

The Port Activation task provides a central console on which the interfaces for any device anywhere in the managed network can be conveniently enabled or disabled. Ports can be taken administratively Up or Down using this task, and all interfaces on a selected device can be activated or deactivated with a series of mouse clicks.

1. From the **Dashboards** tab, select the **Tasks** tab -> **Port Activation**.
2. Choose the **Device Group** from the drop-down list.
3. From the **Device** drop-down list, choose the network device on which port activation will be executed. The **Interfaces** table lists all interfaces on the current device. The **VLAN** and **VLAN Name** columns list the VLAN assigned to each port (VLAN 1/Default resides on all ports without an explicit VLAN assignment). The **OP Status** column will show the current state of each interface.
4. Scroll down the table to locate the interface(s) you want to activate.
5. From the **Admin Status** column, select **Up** (or **Down**) from the drop-down list for the chosen interface.
6. Set any other interfaces on the current device based on your assigned task.
7. Click **Apply**.

The system sends the request to the NetMRI appliance and displays the task configuration sequence.

The Port Activation task will also write the full running configuration to memory, making it the saved configuration. If the user made a change to the running configuration, in parallel with the port activation change, and did not save it, those changes will also be saved.

## Specifying a Port Activation Script

The Port Activation task provides a central console on which the interfaces for any device in the managed network can be conveniently activated. Ports can be taken administratively Up or Down using this task, and all interfaces on a selected device can be activated or deactivated with a series of mouse clicks.

The NetMRI appliance provides the ability to create new automation scripts for many purposes. You may, for example, wish to create a new Port Activation script and use that as an automation task.

To select a different script from the default choice in the software:

1. From the **Dashboards** tab, select the **Tasks** tab. Under the **Port Activation** task, click the settings icon.
2. For Port Activation Options, choose a new script from the **Script Name** drop-down list. The scripts are located on the Trinzi Automation 4000 appliance, and automatically referenced for use by NIOS. By default, the bundled **Port Activation** script is selected.
3. Click **Save**.

The system sends the request to the NetMRI appliance and displays a notification message.

## VLAN Reassignment

VLANs can be reassigned to new interfaces on individual L2/L3 switches in the managed network. A VLAN can have a path across several switches; when you make changes on a given switch, make sure that the path is maintained.

To ensure end-to-end connectivity, you may need to change VLAN port assignments on more than one switch in the path. This feature operates with the VLAN Trunking Protocol (VTP). VLAN switching is changed across one port per switch at a time. Should you need to change VLAN assignments across more than one switch in the path, plan accordingly.

VLANs must already be configured on the switch(es) being changed, and be detected by the NetMRI appliance.

1. From the **Dashboards** tab, select the **Tasks** tab -> **VLAN Reassignment**.
2. Begin by selecting the Device Group from the drop-down list. For **VLAN Reassignments**, you typically choose the Switching device group.
3. From the **Device** drop-down list, choose the switch on which port reassignment will be executed.
4. From the **Port** list, choose the interface to which the VLAN will be reassigned. The **Port** list also shows the Administrative and Operational states of each interface on the current device (Administratively Up/Operationally Down, for example.)  
Note that you can reassign a VLAN to a port that is operationally or administratively Down. The Current VLAN value will show the VLAN to which the selected interface is currently assigned.
5. Choose the new VLAN value for port reassignment from the **New VLAN** drop-down list.
6. Click **Move VLAN**.

The system sends the configuration request to the NetMRI appliance and displays the task configuration sequence.

The VLAN Reassignment task will also write the full running configuration to memory, making it the saved configuration. If the user made a change to the running configuration, in parallel with the port activation change, and did not save it, those changes will also be saved.

## Assigning a New Script to the VLAN Reassignment Task

The NetMRI appliance provides the ability to create new automation scripts for many purposes. You can create and assign a new VLAN Reassignment script and use that for the automation task.

To select a different script from the default choice in the software:

1. From the **Dashboards** tab, select the **Tasks** tab. Under the **VLAN Reassignment** task, click the settings icon.
2. For Port Activation Options, choose a new script from the **Script Name** drop-down list.
3. Click **Save** to commit settings.

The system sends the request to the NetMRI appliance and displays a notification message.

The VLAN Reassignment task will also write the full running configuration to the device's memory, making it the saved configuration. If the user made a change to the running configuration, in parallel with the port activation change, and did not save it, those changes will also be saved.

## Bare Metal Provisioning

The **Bare Metal Provisioning** automated task enables automated installation of new switches and routers into the network. The Trinzic Automation task enables cost and convenience savings by detecting the default behavior of new devices on the network, pointing them to customized TFTP servers from which standardized bare-metal configuration files are downloaded and installed onto the new devices.

The **Bare Metal Provisioning** automated task does not provide NIOS-based optional settings; configuration for this task is done in the Trinzic Automation 4000 NetMRI user interface. The automated task is automatically triggered by detection of a network device requiring configuration.

## Rogue DHCP Server Remediation

All DHCP servers on the network should be under administrative control. If any device offering DHCP leases to clients on the network is not properly administered, it violates many security guidelines and at the very least may cause configuration problems throughout the network. Some events may be unwitting or innocuous (an office worker installing a wireless access point in their cube to share a resource), or may be an attempt to hijack clients and steal information. To prevent such issues, the Rogue DHCP Server Remediation task enables the detection, location and isolation of such devices.

## Viewing Tasks Execution Logs and Approving Tasks in the Task Viewer

You can view the logged results from any task run from the **Automation Tasks** panel through Grid Manager's Task Viewer that displays the following pages:

- **Job History:** Provides a log history of all NetMRI tasks that have recently run, including all automation task types in the dashboard.
- **Issues & Approvals:** Provides links to two important items:
  - **Issue Details:** Displays details about any network issue related to NetMRI tasks and jobs in an Issue Viewer from the NetMRI appliance.
  - **Approve Jobs:** These are jobs that must be approved before the NetMRI appliance can execute the job. For example, the **Isolate Rogue DHCP Server** job must be approved before it will run and attempt to isolate the detected rogue DHCP server in the network.

To view and approve tasks:

1. From the **Dashboards** tab, select the **Tasks** tab.
2. In the **Automation Tasks** panel, click the down arrow icon on the right and select **Task Viewer**.  
The Task Viewer window appears, displaying a scrollable and sortable Job History table. Important columns include the **Start Time**, the **Job ID** (a numeric value with a clickable link to the TAE Job Details Viewer, which will open in a new browser tab), the **Job Name**, the **User** account that executed the task, the job **Status** and the **# Devices** (the number of devices) against which the task ran. The Job History page shows the most recent subset of executed NetMRI jobs. A yellow bar at the top of the table provides a **click here to see more** link, which takes the user to the NetMRI appliance **Job History** page in a new browser tab.
3. If an item appears in the **Issues & Approvals** page, do one of the following in the **Action** column:
  - a. To view an issue in more detail, click an **Issue Details** link. This displays the NetMRI appliance Job Details page in a new browser tab for the selected job.
  - b. To approve a job, click an **Approve Job** link. This displays the **Summary** page of the NetMRI Job Wizard. Click **Approve Job**.
4. Click **Close** to close the Task Viewer.

## Viewing Tasks Execution Logs in NetMRI

You can start NetMRI directly from the Grid Manager's Dashboards tab to view tasks execution logs:

1. From the **Dashboards** tab, select the **Tasks** tab.
2. In the **Automation Tasks** panel, click the down arrow icon on the right and select **Launch NetMRI**.  
NetMRI launches in a new browser tab.
3. Go to **Configuration Management** → **Job Management** side tab → **Scripts** and check the Last Run column.

## Report Clustering Dashboard

The **Reporting Clustering Status** dashboard provides detailed information on the status of the entire indexer cluster. You can get information on the status of each peer node, search head, and indexes. You can view the number of peers (reporting members), searchable copies, and number of copies (buckets).



### Note

The **Reporting Clustering Status** dashboard is available only when you configure the reporting clustering and you must also have the global read-only permissions for Grid Reporting Properties.

The default dashboard includes the following information:

- Indicates whether the reporting data is fully searchable. Displays **Yes** to indicate that all buckets in the cluster have a primary copy.
- Indicates whether the search and replication factors are met.
- Displays the number of indexes and peers that are searchable.

This dashboard might also display the following messages depending on the health of your cluster: Some data is not searchable, All Data is Searchable, Replication factor not met, and Search factor is not met.

## Viewing Reporting Clustering Status

To view the *Reporting Clustering Status* dashboard:

1. From the **Dashboards** -> **Reporting Clustering Status** tab. In the **Peers** tab, you can view the following information:
  - **Peer Name**: The name of the reporting member.
  - **Fully Searchable**: Indicates whether the peer currently has a complete set of primaries and is fully searchable.
  - **Status**: The status of the reporting member.
  - **Buckets**: The number of copies stored on the peer node. The number of buckets for which the peer has copies
- The **Indexes** tab displays the following information:
  - **Index Name**: The name of the indexer.
  - **Fully Searchable**: Indicates whether the peer currently has a complete set of primaries and is fully searchable.
  - **Searchable Data Copies Status**: The status of the reporting member.
  - **Replicated Data Copies**: The replicated number of copies.
  - **Buckets**: The number of buckets for which the peer has copies.
  - **Cumulative Raw Data Size**: The size of the index, excluding hot buckets.
- The **Search Heads** tab displays the following information:
  - **Search Head Name**: Name of the search head.
  - **Status**: The status of the search head.

## Configuring Dashboard Templates

Superusers can specify the tasks an admin group can perform from the **Tasks Dashboard** tab by creating a dashboard template and assigning it to the admin group. When you create a dashboard template, you define the tasks users in an admin group can perform and specify whether the users can configure their own dashboards when they log in to Grid Manager. When you assign a dashboard template to an admin group, all users in this group can see and perform the tasks you define in the template, provided that the users also have the correct

permissions to the objects related to the tasks. For information about administrative permissions, see [Administrative Permissions for Dashboard Tasks](#).

If the assigned template is unlocked, users can configure tasks on their dashboard. If you lock the dashboard template, users cannot configure task packs on their own dashboards.

Superusers can also restrict limited-access users to access only the **Tasks Dashboard** tab when they log in to Grid Manager. These users cannot manage other core network services through Grid Manager. They can only see the **Tasks Dashboard** tab and access only the tasks defined in the dashboard template, if applicable. This feature is useful when you want to define different levels of admin users and restrict them to specific tasks based on their organizational functions. For information about how to set this restriction and creating limited-access admin groups, see [About Admin Groups](#).

To configure and apply dashboard templates, complete the following:

1. Configure dashboard templates, as described in [Configuring Dashboard Templates](#).
2. Assign dashboard templates to admin groups, as described in [About Admin Groups](#).

## Adding Dashboard Templates

Only superusers can configure dashboard templates. Limited-access users may configure task packs depending on the configuration of their assigned dashboard templates.

To add a dashboard template:

1. Log in as a superuser.
2. From the **Dashboards** -> **Tasks** tab, click the Configure icon at the top right corner of a task pack.
3. Select tasks from the Active Tasks table and use the left arrow to move them to the Available Tasks table to hide the tasks, and vice versa. Grid Manager displays the tasks you place in the Active Tasks table. Repeat the steps for all task packs.
4. At the top right corner of the Tasks Dashboard panel, click the Configure icon -> **Configure Template**.
5. In the Dashboard template configuration section, click **Create new template**.
6. In the *Save Dashboard Template* dialog box, complete the following:
  - **Name**: Enter a name for the new dashboard template.
  - **Locked**: When you select this checkbox and assign this template to an admin group, users in the admin group can only perform the tasks you configure to appear in this template. They cannot configure their dashboards. When you clear this checkbox, users can still only see the tasks you configure for this template, but they can now configure tasks in the task packs on their dashboards. Note that when you lock a template, it applies to all users in the admin group, including those who have customized dashboards.
7. Click **Save & Close**.

The appliance saves the template and adds it to the **Template** drop-down list.

## Resetting Dashboard Templates

Only users with an unlocked dashboard template assigned can reset their dashboards to the template that was originally assigned to them. Users with locked dashboard template cannot configure or reset their dashboards. Also, only superusers can configure dashboard templates.

To reset a dashboard template:

1. Select the **Dashboards** -> **Tasks** tab.
2. For superusers: At the top right corner of the Tasks Dashboard panel, click the Configure icon -> **Reset**. Note that the Configure icon appears only if you are a superuser.  
For limited-access users: At the top right corner of the Tasks Dashboard panel, click **Reset**.  
The appliance reset your dashboard to the original dashboard template that was assigned to your admin group.

## Modifying Dashboard Templates

You can modify an existing dashboard template by locking or unlocking it, and adding or removing tasks from a task pack. However, you cannot change the name of the template. When you change the name of a template, the



appliance clones the template and adds the new template to the list. Note that when you modify a locked template that is assigned to an admin group, users in the group automatically adopt the changes you make to the template the next time they log in to Grid Manager.

To modify a dashboard template:

1. From the **Dashboards** -> **Tasks** tab, click the Configure icon at the top right corner of the panel.
2. In the Dashboard template section, select the template you want to modify from the **Template** drop-down list. Note that Grid Manager displays [L] before the name of a locked template.
3. In the task pack, click the Configure icon at the top right corner.
4. Select tasks from the Active Tasks table and use the left arrow to move them to the Available Tasks table to hide the tasks, and vice versa. Grid Manager displays the tasks you place in the Active Tasks table. Repeat the steps for all task packs.
5. Click **Save**.
6. In the *Save Dashboard Template* dialog box, modify other information, as described in [Configuring Dashboard Templates](#).
7. Click **Save & Close**.

## Deleting Dashboard Templates

Only superusers can delete dashboard templates. To delete a dashboard template that is currently assigned to an admin group, you must first unassign the template from the admin group. For more information about creating limited-access admin groups, see [About Admin Groups](#).

To delete a dashboard template:

1. From the **Dashboards** -> **Tasks** tab, click the Configure icon at the top right corner of the panel.
2. In the Dashboard template section, select the template you want to delete from the **Template** drop-down list.
3. Click **Delete**.
4. In the *Delete Dashboard Template* dialog box, click **Yes**.

## Assigning Dashboard Templates

After you create a dashboard template, you can assign it to an admin group. Admin users in this admin group can access the tasks you define in the template.

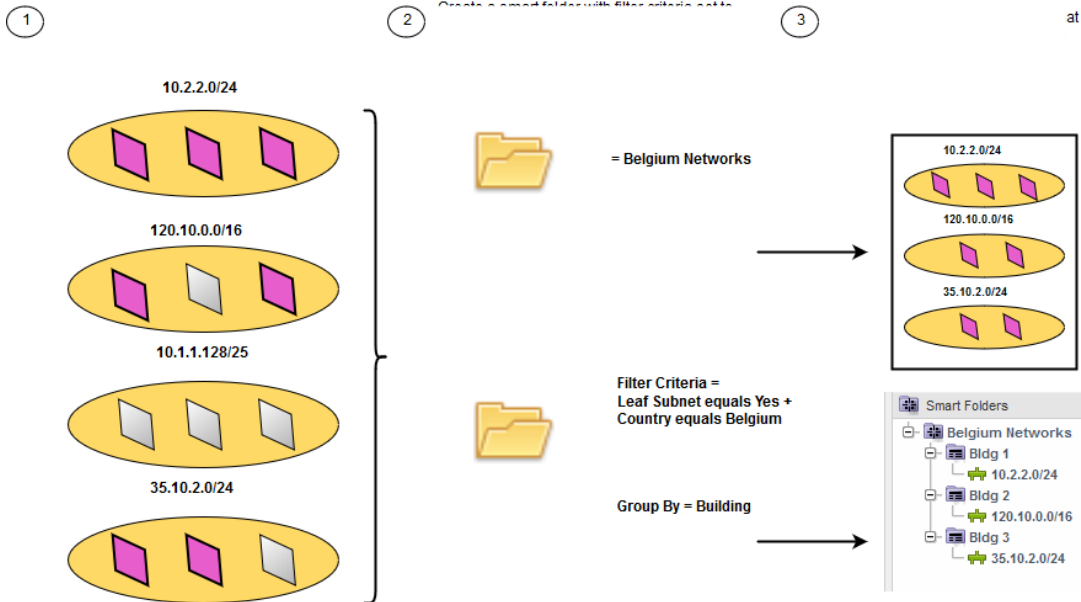
To assign a dashboard template to an admin group, see [About Admin Groups](#).

## Smart Folders

Use smart folders to organize your core network services data. Depending on your administrative roles and business needs, you can filter your data by object types, names, extensible attributes, and discovered data such as conflicts, unmanaged data, or the virtual entity data, and then place the filtered results in a smart folder. You can also group the filtered results by defining up to 10 extensible attributes as the Group By rules. For example, you can create a smart folder that contains all the networks you manage in Belgium, and then group the networks by building number, as illustrated in the figure below.

Once you set up a smart folder, the appliance displays up-to-date information based on your filter and grouping criteria each time you access the folder. You can also view and modify object information in the folder. For information, see [Viewing and Modifying Data in Smart Folders](#). With smart folders, you can organize your data in a meaningful way and quickly obtain the information you need to perform specific tasks without searching the entire database.

## 1 Creating Smart Folders



Before you set up your smart folders, decide how you want to organize your data. You can specify search and Group By criteria to help you group information in a meaningful way. You can also decide whether you want to include objects that do not contain attribute values when you use the Group By criteria to group filtered data by extensible attributes. For information, see [Creating Smart Folders](#). Note that a smart folder becomes invalid when you delete an extensible attribute that the folder uses as a filter or Group By criterion. You must redefine the extensible attribute and reconfigure the folder criteria to validate the smart folder.

In Grid Manager, you can create smart folders in both the Global Smart Folders and My Smart Folders panels. In Global Smart Folders, you can create smart folders to which other administrators can create links. Only administrators with superuser accounts can create, edit, and delete global smart folders. You can create personal folders as well as links to global smart folders in My Smart Folders. For information, see [My Smart Folders and Predefined Smart Folders](#). Each smart folder you create can contain up to 2,000 objects. When the number of objects exceeds 2,000, Grid Manager sorts and displays the first 2,000 objects only. It also displays a warning message at the top of the panel. In this case, you may want to redefine your filter criteria to further refine the filtered data in your smart folders.

To create smart folders, follow these procedures:

1. Determine how you want to organize your core network services data.
2. Identify the fields that you want to use to group networks or define extensible attributes for the data that you want to track. For information about extensible attributes, see [Managing Extensible Attributes](#).
3. Create smart folders in either the My Smart Folders or Global Smart Folders panel. For information, see [Creating Smart Folders](#).



### Note

Infoblox strongly recommends that you use **Type** as one of the filter criteria to improve system performance.

## Global Smart Folders

You can create global smart folders to share among administrators. You must log in as a superuser account to create, edit, and delete global smart folders. All other users have read-only access to global smart folders. You can create as many folders as you need in Global Smart Folders, but Grid Manager displays a maximum of 500 smart folders in the list panel. You can also save a local copy of an existing folder, depending on your administrative permissions. For information, see [Saving Smart Folders Data](#).

Grid Manager displays a list of global smart folders in the list panel.

When you log in as a superuser and mouse over a global smart folder, the following icons appear:

- **Information:** Displays information about the selected smart folder. Information includes comments and filter criteria for the folder. It also displays the Group By rules.
- **Edit:** Click this icon to edit the definition and filter criteria for the smart folder.
- **Create link:** Click this icon to create a link to the smart folder. The link to this folder is placed in My Smart Folders.
- **Delete:** Click this icon to delete the smart folder. This operation does not affect the objects that are in the folder. Only the smart folder is deleted.

## Creating Smart Folders

You can create personal smart folders in My Smart Folders. You can also create global folders to share among administrators in Global Smart Folders when you log in as a superuser account. Each time you access a smart folder, you obtain up-to-date information about the core network services data that match the filter criteria you set for the folder. You can also set Group By rules to group the filtered data by extensible attributes. Grid Manager displays a hierarchical view of the data using the Group By rules you define.

To create a smart folder:

1. Click the **Smart Folders** tab.
2. Click the **My Smart Folders** tab to create a personal smart folder.  
or  
If you logged in with a superuser account, click the **Global Smart Folders** tab to create a global smart folder.
3. Click **Create**.
4. In the *SmartFolder* data panel, complete the following:
  - **Name:** Enter the name of the smart folder.
  - **Comment:** Optionally, enter additional information about the smart folder.
  - In the first drop-down list, select a field as the filter. You can select a network view or a record type as the filter. Grid Manager highlights extensible attributes and Active Directory Sites in gray. You can also group the default data by adding a Group By rule without adding a filter.  
Note that Infoblox strongly recommends that you use Type as the first criterion to improve system performance.
  - In the second drop-down list, select an operator for the filter.
  - Enter or select a value for the selected field and operator. Depending on the field and operator that you select, the field can be a text or an integer field. It can also be a drop-down list or a calendar widget. For example, if you select **Network/Zone/Range/Member** in the **Type** field, Grid Manager displays all the networks, zones, DHCP ranges, and members in the results table. The results table may display the name in its native characters if the name was originally entered as an IDN (Internationalized Domain Name). For example, the **Name** column in the results table displays 网络, a zone name in Chinese.
  - Optionally, click **+** to add another filter. You can also click **Apply** to view the filtered data in the results table.
  - Optionally, select the **Group Results** checkbox to organize the filtered data. You can also disable a **Group By** filter by deselecting the checkbox.
  - From the Group by drop-down list, select an extensible attribute or an Active Directory Site by which you want to group the filtered data. For example, if you want to group the filtered data by building number, you can select **Building** from the drop-down list. To add additional Group By rules, click the **+** icon, and then select a field from the drop-down list. You can apply up to 10 Group By rules. You can also delete a rule by selecting the rule and clicking the **-** icon.
  - After you add all filter criteria and **Group By** rules, click **Apply**. Grid Manager displays the filtered data in the results table. Note that in the **Name** field, the appliance highlights disabled DHCP objects in gray. A DHCP object can be a DHCP range, fixed address, reservation, host configured for DHCP, or roaming host with an allocated IP address. If you select to include objects with no attribute values in the **Group By** rules, the appliance may take longer to process the results. You can also choose to display Active Directory Site values in the table.
5. Click **Save** to save the smart folder.

## Viewing and Modifying Data in Smart Folders

After you set up a smart folder, the appliance searches for matching objects based on the filter criteria you specified for the folder. Grid Manager also groups the objects by the Group By rules you specify. If you select to include objects with no attribute values, the appliance may take longer to process the results. Each smart folder you create can contain up to 2,000 objects. When the number of objects exceeds 2,000, Grid Manager sorts and displays the first 2,000 objects and shows a message at the top of the panel. In this case, you may want to redefine your filter criteria to further refine the filtered data in your smart folders.

Grid Manager displays smart folders hierarchically in a tree view based on your Group By rules as follows:

- Smart Folder section in the Finder panel.
- Selectors from which you can select a smart folder.

In the smart folder list panel however, Grid Manager displays all the smart folders in a flat list. You can modify some of the data in the table. Double-click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [About the Grid Manager Interface](#).

In the smart folder data panel, Grid Manager displays the first hierarchical level of the smart folder based on your Group By rules. If you do not configure any Group By rules, Grid Manager displays all the objects in the results table. If you select to include objects with no attribute values, Grid Manager also includes these objects in the hierarchical view. Depending on your Group By rules, you can view detailed information about the objects by clicking the object link and drilling down to the lowest hierarchical level, and then opening an object. To go back to a previous hierarchical view, click the link of the corresponding level in the breadcrumb.

To view detailed information about an object, complete the following:

1. In the Smart Folder data panel, click the object link until you drill down to the last hierarchical level of the folder.
2. Grid Manager displays the following information:
  - **Name:** The name or IP address of the object.
  - **Comment:** Information about the object.
  - **Type:** The object type.
  - **Site:** The site to which the object belongs. This is one of the predefined extensible attributes. You can also select other available extensible attributes for display, and sort the data in ascending or descending order by column.
3. Select an object checkbox, and then perform one of the following:
  - Click the Open icon to display the data in the network list or IP address list.
  - Click the Edit icon to modify or schedule the modification of the object configuration. Grid Manager displays the corresponding editor depending on the object you select.
  - Click the Delete icon to delete the object or click the Schedule Deletion icon to schedule the deletion of the object.

You can also print or export the data in this panel. For information, see [Using the Grid Manager Interface](#).

## Modifying and Deleting Smart Folders

After you create a smart folder, you can either modify its filter and grouping criteria. You can even delete personal smart folders in My Smart Folders.

### Modifying Smart Folders

1. Go to **Smart Folders**.
2. Click **My Smart Folders** to modify personal smart folders.  
or  
Click **Global Smart Folders** to modify global smart folders if you logged in with a superuser account.
3. Mouse over to the smart folder that you want to modify.

4. Click the Edit icon. You can also click the Edit icon next to the name of the smart folder in the data panel.
5. Make the appropriate changes in the Smart Folder data panel as described in [Creating Smart Folders](#).

## Deleting Smart Folders

You can delete personal smart folders in My Smart Folders. However, you must log in as a superuser account to delete global smart folders.

To delete a smart folder:

1. Click the **Smart Folders** tab.
2. Click the **My Smart Folders** tab to delete personal smart folders.  
or  
Click the **Global Smart Folders** tab to delete global smart folders.
3. Mouse over to the smart folder that you want to delete.
4. Click the Delete icon. In the *Delete Smart Folder* dialog box, click **Yes**.

## Saving Smart Folders Data

You can make a copy of an existing smart folder, add or change filter criteria, and then rename the folder accordingly. You can also create a local copy of the global smart folder in its current state for editing purposes. In My Smart Folders, you can save a folder copy only in My Smart Folders. In Global Smart Folders however, you can save a folder copy in either My Smart Folders or Global Smart Folders. You must have superuser permissions to save a global smart folder copy in Global Smart Folders. Note that when the original global smart folder is updated, information in your local copy is not updated.

To save a copy of a smart folder:

1. Click **My Smart Folders** to save a folder copy in this tab.  
or  
Click **Global Smart Folders** to save a folder copy in either this tab or My Smart Folders. To save a smart folder copy in Global Smart Folders, log in as a superuser account.
2. Select the smart folder that you want to save as a copy.
3. Click **Save copy as**.
4. Grid Manager saves the folder copy in My Smart Folders when you save the folder copy in this tab.  
or  
The *Save Smart Folder As* dialog box appears when you perform this function in Global Smart Folders. Select one of the following:
  - **MySmartFolders**: Saves the copy in My Smart Folders.
  - **Global Smart Folders**: Saves the copy in Global Smart Folders.Click **OK**.



### Note

For the folder copy, the appliance appends the word Copy to the original name of the smart folder. You can change the name of the folder at anytime by editing the folder.

## Printing and Exporting Data in Smart Folders

You can print the list of networks that are on the current smart folder page, or you can export all the data in CSV (comma separated value) format. For more information, see [Scheduling Tasks](#) and [Importing and Exporting Data using CSV Import](#).

## My Smart Folders and Predefined Smart Folders

Grid Manager has two types of smart folders available:

- My Smart Folders
- Pre-defined Smart Folders

### My Smart Folders

In My Smart Folders, you can create personal smart folders and links to global smart folders. When you create links to global smart folders, you can only view information in the folders. However, you can create a local copy of the global smart folder in its current state for editing purposes. Note that when the original global smart folder is updated, information in your local copy is not updated. For information, see [Saving Smart Folders Data](#). When you delete a link to a global smart folder in this tab, only the link is deleted. There is no impact on the information in the original global smart folder.

Grid Manager displays a list of smart folders in the list panel. The same list of smart folders is also displayed in the *Finder* panel. For more information, see [About the Grid Manager Interface](#).

When you mouse over a smart folder in the list panel, the following icons appear:

- **Information:** Displays information about the selected smart folder. Information includes comments and filter criteria of the folder. It also displays how you grouped the filtered data.
- **Edit:** Click this icon to edit the definition and filter criteria for the smart folder.
- **Delete:** Click this icon to delete the smart folder. This operation does not affect the objects or networks that are in the folder. Only the smart folder is deleted.

### Predefined Smart Folders

The appliance can detect remote clients through their DHCP fingerprints, or through device information discovered through SNMP and other device and network detection protocols. You can use predefined smart folders to view lease history, IP addresses, network infrastructure devices, and other related information for remote clients that contain DHCP fingerprint information related to the following device groups:

- **Smartphone,PDA,Tablet Devices:** Includes all devices that were detected as smartphones, PDAs, and tablets.
- **Microsoft Windows Devices:** Includes all devices that were detected to be running Windows OS.
- **Apple MAC OS Devices:** Includes devices that were detected to be running MAC OS.
- **Conflicts:** Includes hosts detected in the network that have a MAC Address conflict. A discovered host has a MAC address conflict when its MAC address is different from that specified in its fixed address, DHCP lease, or host record.
- **Discovered Routers/Switches:** Includes core network infrastructure devices of the specific Router, Switch, or Switch-Router types discovered by NIOS using the discovery feature set. Clicking on a device name opens the device page under **DataManagement->Devices** and shows the **Interfaces** page for the chosen device.
- **Active Directory Sites:** Includes all names that were detected as Active Directory Sites. Clicking on an Active Directory Site opens the *Active Directory Site Properties* editor where you can edit the name, add, or delete networks that are associated with the Active Directory Site. For more information, see [Managing Active Directory Sites and Associated Networks](#).
- **Gaming Console Devices:** Includes devices that were detected as gaming consoles.
- **Router and Wireless Access Point Devices:** Includes devices there were detected as routers or wireless access point devices.
- **Unmanaged:** Shows all unmanaged devices.

### Related topics

[DHCP Fingerprint](#)

[Infoblox Network Insight](#)

# Importing and Exporting Data using CSV Import

Use **CSV Import** to import DNS, DHCP, IPAM data, and subscriber site data through Grid Manager. You can use this feature to migrate or add new data, overwrite existing data, merge new data with existing data, delete existing data, or replace certain existing data in the database.

To import new data, you must first prepare a data file (include all required fields and follow the proper syntax), and then start an import through Grid Manager. You can also export existing data to a data file, modify the data, and then import the modified data to the database. You can either overwrite existing data with the modified data or merge new data with the existing data. You can also delete data that is no longer required or replace certain existing data with new data in the file. Note that the replace option is valid for authoritative zone data only whereas other options are valid for all supported objects including zones. The replace operation creates a snapshot or a backup of the existing data in the database before replacing the database with the data in the imported CSV file.

The appliance supports CSV import for most record types. You can use IDNs and punycode for the domain name field for most of the DNS object types. For information about IDNs and punycode, see [Support for Internationalized Domain Names](#). Only superusers can import A and AAAA records with a blank name. Limited-access users must have read/write permission to **Adding a blank A/AAAA record** in order to import A and AAAA records with a blank name, otherwise the CSV import operation might fail. You can assign global permission for specific admin groups and roles to allow to import A and AAAA records with a blank name. For more information, see [Administrative Permissions](#). For each supported record type, you must include all required fields in the header row of the dataset that you want to import. For a list of supported record types and specific guidelines for creating a data file, refer to the [Infoblox CSV Import Reference](#).

To import a data file:

1. Create a data file if you do not already have one. Follow the guidelines for the supported objects to ensure that you include all the required fields in the file. For more information, refer to [Infoblox CSV Import Reference](#). You can also export existing data and then update the file for re-import. For more information, see [Exporting Data to Files](#) below.
2. Configure import options, see below.
3. Start a CSV import. For information, see [Configuring Import Options](#) below.



## Warning

CSV imports and operations that involve massive data such as deleting large zones and recursive deletion of networks and all child objects, will significantly affect member performance, resulting in service outage.

When you submit multiple CSV imports, the appliance puts the import jobs in queue and executes them one at a time in the order they are submitted. When a job is being executed, it is in the **Import in progress** state. When a job is in queue for execution, it is in the **Import pending** state. You can import multiple CSV files at a time, but at any given time you can execute only one single task. Note that only one task at a time will be in the **Import in progress** state, while the others are in the **Import pending** state. You can view the status of each import job through **CSV Job Manager**. Superusers can view all import jobs while limited-access users can view only the jobs they submitted.

To access **CSV Job Manager**, from the **Data Management** tab, click **CSV Job Manager** from the Toolbar and select **Jobs Manager**, or from the **Tasks Dashboard**, click **CSV Import** in the IPAM Task Pack. Superusers and limited-access users that have applicable configurations and permissions can perform CSV imports and exports. For information about user permissions for CSV imports and exports, see [CSV Import User Permissions](#) below.

You can do the following in **CSV Import**:

- Add, overwrite, append, replace or delete data through the imported CSV file, as described in [Configuring Import Options](#).
- Verify the content in the CSV file, as described in [Configuring Import Options](#).
- View a list of CSV import jobs, as described in [Creating a Data File for Import](#).
- Add and start CSV import jobs, upload data files, stop CSV imports, or edit the options of the uploaded file, as described in [Modifying CSV Import Jobs](#).
- Delete uploaded jobs, as described in [Deleting Uploaded Jobs](#).

- Download the following: imported files, import errors, import results, or snapshots, as described in [Downloading Files](#).
- Select a pending or saved job, and then click the Cancel icon to cancel the job.
- Click the Refresh icon to refresh the **CSV Job Manager**.



## Notes

- The list of CSV import jobs is not restored when you restore a backup file or when you promote a master candidate.
- Superusers can view any jobs in the **CSV Job Manager**, and limited-access users can only view jobs they submitted.
- A non-superuser can import or export CSV files for subscriber records.

## CSV Import User Permissions

Superusers can perform any CSV import tasks. You must assign limited-access users the correct configurations and permissions so they can perform CSV imports and exports. For information about how to configure the CSV Import task for limited-access users, see [Configuring Dashboard Templates](#). Limited-access users can import data to which they have proper permissions. For information about admin permissions, see [About Administrative Permissions](#). Changes you make to user permissions can affect CSV import and export behaviors. The following table lists actions performed on user permissions and the corresponding effects on CSV imports and exports:

Actions taken on user permissions	CSV import and export behaviors associated with the affected user account
Delete a user account	<ul style="list-style-type: none"> <li>• CSV import jobs remain in the system and are accessible by superusers only.</li> <li>• All pending import jobs cannot be executed due to authentication failures.</li> <li>• If the action is taken during an import job that is in progress, the rest of the import will fail.</li> <li>• All stopped and successful jobs are available to superusers only.</li> </ul>
Modify a user account	<ul style="list-style-type: none"> <li>• If the action is taken during an Import in progress import job, the rest of the import will fail.</li> </ul>
Remove user permissions in a user account	<ul style="list-style-type: none"> <li>• Pending import jobs may be completed with errors.</li> </ul>

## CSV Import Limitations

Ensure that you understand the following limitations before you start an import:

- You can import multiple CSV files at a time, but at any given time you can execute only one single task. The import tasks are queued. Note that only one task at a time will be in the **Import in progress** state, while the others are in the **Import pending** state.
- Do not use UTF-8 characters in the CSV file name.
- When you perform a CSV import that includes objects that have scheduled changes or updates associated with them, the import fails. Only superusers can cancel the scheduled changes.
- When you stop an import, the appliance completes the import of the data row it is currently processing before it stops the import. You cannot resume the import from where it stopped.
- You cannot roll back to previous data.
- The following data cannot be imported: Microsoft management, DNSSEC, and GSS-TSIG data.



- CSV import does not support DNSSEC zones, though resource records added for a signed zone are supported.
- Only editable data can be imported. Discovered data cannot be imported or manipulated.
- When you promote a new Grid Master during an import, the import stops; and it does not restart on the new Grid Master. When a failover occurs during an import, the import stops on the old active node, and it does not restart on the new active node.
- It may take longer than expected to import a large number of DHCP ranges that are associated with a single MAC address filter.
- When a CSV import starts, the appliance validates the first 100,000 rows of data in the CSV file. If the file contains more than 100,000 rows of data, the appliance validates the rest of the data as the import progresses.
- The appliance supports up to one million rows of data in each CSV import.
- You cannot import network containers.
- To successfully import RIR (Regional Internet Registries) organizations, you must also specify the maintainer password. Note that the password field is not exported during a CSV export. For information about RIR updates, see [RIR RIR Registration Updates Registration RIR Registration Updates Updates](#).
- You can use the **Replace** operation to replace the current data in the database with the data in the imported CSV file. Note that the replace option is valid for authoritative zone data only whereas the other options are valid for all supported objects including zones. For more information, see [Configuring Import Options](#) below.
- The **Replace** operation is available only for authoritative zones. This operation does not support DNS records that are automatically generated or exported, but it supports NS records that are created manually.
- Use the delete function to delete import jobs that are uploaded. You can delete the content of a CSV file that you have imported to the database. Note that you cannot delete jobs that are already imported.
- When you import CSV files for NS record updates, you must specify a value for **zone\_nameservers**. NIOS displays an error message if you do not specify a value for this field when you import the CSV file.
- When you perform a CSV export of automatically created NS records using Infoblox API, the **zone\_nameservers** field will have an empty value. Therefore, if you import the previously exported CSV file that includes automatically created NS records through the Infoblox GUI, then the CSV import fails, and Grid Manager displays an error message.
- If you upload a file and preview the file using the **Preview** option, and later update the content of the same CSV file, and then try to view the edited file using the same *Preview* wizard, you may not be able to see the changes. Infoblox recommends that you start a fresh CSV import to upload the edited file and navigate to the *Preview* wizard to repreview the file.
- You cannot perform the CSV import operation on a Microsoft Server zone object, but NIOS allows you to perform the CSV import operation on records within a Microsoft Server zone. You may not see an error message when you perform a CSV import using the replace operation on a Microsoft Server zone.
- You cannot perform merge, custom, and replace operations for subscriber records.
- A non-superuser can import or export CSV files for subscriber records.

## Configuring Import Options

You can import CSV files and perform various operations to update the data in the database. You can choose from several import options: add, override, merge, delete, and replace. You can add new rows from the imported file to the database, overwrite existing rows in the database, append rows to the existing rows in the database, delete existing rows in the database, or replace the existing rows in the database. You can verify whether the data in the imported file is appropriate using the **Test** option before you import the file to the database. You can also view the results and progress details of the operation.

To import a CSV file, complete the following:

1. From the **Data Management** tab, click **CSV Job Manager** from the Toolbar.
2. In the *CSV Job Manager* wizard, select **CSV Import** and click the New CSV import job icon.
3. In the *New CSV Import Wizard*, complete the following:
  - **Type of Import**  
**For all supported objects (including zones)**
    - **Add:** Select this to add new rows from the imported CSV file to the database. NIOS updates the database with the new data that you have added to the imported CSV file and retains the rows that do not have any changes.
    - **Override:** Select this to overwrite the existing data in the database with the data from the uploaded file. You cannot add new rows or delete existing rows. If you want to overwrite values in the

required fields, you must include the required fields and the corresponding `_NEW_XXXX` fields in the data file.

- **Merge:** Select this to add values from the imported CSV file to the existing columns in the database that do not have any data. It does not overwrite the existing data, even if the data file contains new values for certain fields. If you want to overwrite values in the required fields, you must include the required fields and the corresponding `_NEW_XXXX` fields in the data file.
  - **Delete:** Select this to delete the rows in the imported CSV file from the database.
  - **Custom:** Select this to apply custom import actions for individual data rows in your CSV file. When preparing the CSV file for import with the **Custom** option, add an **IMPORT-ACTION** column to the file and specify a custom import action for each data row. Use the following abbreviations for import actions: 'I' (INSERT), 'M' (MERGE), 'O' (OVERRIDE), 'IM' (INSERT+MERGE), 'IO' (INSERT+OVERRIDE), 'D' (DELETE).
- **For zones only**
    - **Replace:** Select this to replace the contents of the database with data provided in the CSV file. NIOS cancels the replace operation and will not save the changes if it encounters an error. You can replace the DNS records of a zone by importing a zone file that was exported previously. You can only replace DNS records that are manually created. NIOS generates a results file listing the file name, action performed, date and time, and result at the end of the validation. You can view the results file only after the replace operation is complete. NIOS generates the backup file automatically for every replace operation and saves it in the Infoblox Grid. Note that the CSV file must contain data for one authoritative zone only, that is, you cannot insert records from different authoritative zones into a single CSV file for replace operation.  
Note: The replace operation might affect system performance if you try to replace a zone with a lot of changes. Infoblox recommends that you perform the replace operation for large import files (more than 10,000 rows of changes) during non-peak hours. This operation ignores `_new_XXX` fields in the imported CSV files.
4. Click **Next** to import the CSV file.
    - **Import Type:** Displays the type of import option you have selected.
    - Click **Choose** and select the CSV file that you want to import, and then click **Open**.
    - **On Error:** Select one of the following to tell the appliance what to do when it encounters an error during an import:
      - **Stop import:** Select this to stop the data import once it encounters an error in the uploaded file. Grid Manager displays the row number at which it stops the import when it encounters an error. NIOS saves the changes made to the CSV file before an error occurs. For example, if there are 100 rows of data and you select this option, and there is an error in row 90, the appliance displays **90 of 100 completed, 1 error**.
      - **Skip to the next row and continue:** Select this to skip over errors and continue the data import. You can download an error report to identify the erroneous data. NIOS displays the total number of rows it has processed by skipping over. For example, if there are 100 rows of data and you select this option, the appliance displays **100 of 100 completed, 1 error**.
  5. Click **Next** to preview your CSV file. In the File Preview table, Grid Manager displays the header row, the first six rows, and up to 15 columns of the imported data. You cannot edit the data here. Field names with asterisks ★ indicate required fields. Note that you must define these fields in the imported file. If any of the required fields are missing, the appliance generates an error during the import operation. You can do the following in this wizard:
    - **Import type:** The type of import option you have selected.
    - **Filename:** The name of the CSV file you have selected.
    - **Separator:** Select a separator for your CSV file from the drop-down list. The default value is **Comma**.
    - **OnError:** The option you have selected.
  6. Click **Test** to verify the content in your CSV file. Click **Yes** in the *Test CSV Import for Replace Operation* dialog box to verify the content or click **No** to cancel the operation. NIOS automatically analyzes the data in the imported file for any syntax errors or other violations. You can also view a detailed report of the file that you are importing. Note that you can run the test as a background task. This report also displays information about the number of deleted, updated and added files. It also displays error messages, if any. NIOS generates a results file listing the file name, action performed, date and time, result, and the number of failures at the end of the validation. You can view the results file only after the replace operation is complete.  
Note that the **Test** button is enabled only when you select the **Replace** operation and is disabled for other import options.

7. Click **Import** to import the CSV file to the database. Click **Yes** in the dialog box to import the CSV file or **No** to cancel the operation.
8. You can view the progress and results of your import operation in the *CSV Import Progress* wizard. This wizard displays the following information:
  - **Import type:** The type of import option you have selected.
  - **Filename:** The name of the CSV file you have selected.
  - **Separator:** The separator you have selected for your CSV file. The default value is **Comma**.
  - **On Error:** The option you have selected when the import operation encounters an error.
  - **Current status:** If an import is in progress, this field displays its current status. Otherwise, it displays the date and time of the last import.
  - **Last action:** Displays the last operation and the admin who initiated it.
  - **Rows Completed:** The number of rows of data the import has processed. Depending on the import options, Grid Manager displays either the row number at which it stops an import when it encounters an error or the total number of rows it has processed by skipping over the erroneous data. For example, if there are 100 rows of data and you select "On error: Stop importing," and there is an error in row 90, Grid Manager displays **90 of 100** here. If you select "On error: Skip to the next row and continue," Grid Manager displays **100 of 100** here and displays **1 in Rows with Errors**.
  - **Rows with Errors:** The number of rows of data the import has detected errors. Click **Download Errors** to download the CSV file that contains the fields and the rows of erroneous data. You can use this report as a reference to update the data file before you import the file again.
9. You can also perform the following actions if needed:
  - To cancel the import operation, click **Stop Import** before the operation is complete.
  - To close the wizard and execute the operation in the background, click **Close & Run in Background**.
  - When the operation is complete, you can click **Download errors** to download and view the errors. The **Download errors** button is enabled only if the operation encounters errors.
  - Click **Save & Close** to save the operation and close the wizard.



#### Note

Superusers can view all CSV import jobs and limited-access users can view only the jobs they submitted.

## Viewing CSV Import Jobs

You can view the status of import jobs. To view the status:

1. From the **Data Management** tab, click **CSV Job Manager** from the Toolbar.
2. In the **CSV Job Manager** wizard, click **CSV Import**. Grid Manager displays the following information about the import jobs that were submitted in the past 30 days:
  - **User Name:** The admin user who submitted the CSV import. Only superusers can view this column.
  - **Status:** The current status of the import job. The status can be one of the following:
    - **Import successful:** The import is completed without errors. Check the **Message** field for information about the import.
    - **Import unsuccessful:** The import is completed, but with errors. Check the **Message** field for information about the error message.
    - **Import pending:** The job is in queue for execution.
    - **Import inprogress:** The job is being executed.
    - **Import stopped:** The job has been stopped. You can select the job and restart the import.
    - **Test successful:** Test is completed without errors. Check the **Message** field for information about the test.
    - **Test unsuccessful:** Test is completed, but with errors. Check the **Message** field for information about the error message.
    - **Test pending:** Test is in queue for execution.
    - **Test inprogress:** Test is in progress.
    - **Test stopped:** Test has been stopped. You can select the job and restart the import.

- **Saved file:** The data file has been uploaded, but the import has not started. Note that after a product restart, which can be caused by a failover, all **Import in progress** jobs go into **Import stopped** state; all **Import pending** jobs continue to be queued for execution.
- **Submitted:** The timestamp when the job was submitted.
- **Completed:** The timestamp when the job was completed. This field is blank if the job has not been completed yet.
- **File Name:** The CSV data file name.
- **Message:** This field displays the number of rows of data that have been processed and the number of rows of data the import has detected errors. Depending on the import options, Grid Manager displays the row number at which it stops the import when it encounters an error, or the total number of rows it has processed by skipping over the erroneous data. For example, if there are 100 rows of data and you select "On error: Stop importing," and there is an error in row 90, the appliance displays 90 of 100 completed, 1 error. If you select "On error: Skip to the next row and continue," the appliance displays **100 of 100 completed, 1 error**.
- **FileSize:** The size of the imported CSV file.

Superusers can view all CSV import jobs and limited-access users can view only the jobs they submitted.

## Modifying CSV Import Jobs

You can modify the options of the CSV file that you have already uploaded, delete the jobs that are uploaded, or download the uploaded file or error file. After you configure the import options, you can select a data file and start an import operation or upload a data file. For more information about configuring import options, see as described in Configuring Import Options.

To edit the options of a file, complete the following:

- From the **Data Management** tab, click **CSV Job Manager** from the Toolbar.
- In the *CSV Job Manager* wizard, select **CSV Import** and select the import job that you want to update, click the Action icon  
  
and select **Edit**.
- In **Edit-CSV Import Job**, select a type of import and perform the operations mentioned in Configuring Import Options.
- Click **Download** to download the uploaded file, snapshot file, or the results file.
- Click **Save & Close** to save the changes.

## Deleting Uploaded Jobs

You can delete import jobs that are uploaded. You cannot delete jobs that are already imported. You can delete the content of a CSV file that you have imported to the database. Note that the CSV import files and the backed up files are saved for a period of 30 days, but the size limit is set to 1 GB. If these files increase in size, NIOS removes the older files from the Grid. NIOS generates a syslog message if it encounters an error when generating the backup file.



### Note

When you delete a parent object from the CSV file, the child objects associated with the parent objects are also deleted.

To delete uploaded jobs, complete the following:

1. From the **Data Management** tab, click **CSV Job Manager** from the Toolbar.
2. In the *CSV Job Manager* wizard, select **CSV Import** and select the import job that you want to delete, click the Action icon and select **Delete** or click the Delete pending job icon.
3. Click **Yes** to delete the uploaded job or **No** to cancel the operation in the *Cancel Import Job* wizard.

## Downloading Files

You can download various types of files based on the import operation that you have selected. You can download the following files: uploaded, error, results, and snapshot. Superusers can download the original imported file.

### Downloading Uploaded or Error Files

You can download CSV files that are already uploaded or download error files to check the errors that the import operation encountered. The download options are valid for all import operations, except replace.

To download the file, complete the following:

- From the **Data Management** tab, click **CSV Job Manager** from the Toolbar.
- In the *CSV Job Manager* wizard, select **CSV Import** and select the import job that you want to download, click the Action icon and select **Download**.
  - **Uploaded File**: Select this to download the uploaded CSV import file.
  - **Error File**: Select this to download the error file. This option is enabled only if the import operation encountered an error.

You can export these files to your local system.

### Downloading Uploaded, Snapshot, or Results Files

You can view the uploaded, result, and snapshot files. The snapshot and results files are enabled only for replace operation.

To download the file, complete the following:

1. From the **Data Management** tab, click **CSV Job Manager** from the Toolbar.
2. In the *CSV Job Manager* wizard, select **CSV Import** and select the import job that you want to download, click the Action icon and select **Edit**.
3. In **Edit - CSV Import Job**, click the arrow beside the **Download** option and select one of the following:
  - **Uploaded File**: Select this to download the uploaded CSV import file. The uploaded report displays the content of the imported file, before it was uploaded, that is before the content of the database is changed.
  - **Snapshot File**: Select this to download the backup file. NIOS automatically creates a backup of the database before replacing the content of the database with the content in the imported CSV file. This option is enabled only for the replace operation. NIOS generates an error message and saves it in the syslog and the infoblox.log file when the backup file is not generated. You can download the backup file after the replace operation is complete. NIOS saves the backup file and other results file in the Grid for a period of 30 days. The name of the backup file has the following format:  
`csv-snapshot-10-[view\]-[zone\]-[timestamp\].csv` where 10 is the import ID of the CSV import task, view is the DNS view name, zone is the FQDN of the zone being replaced and timestamp is the timestamp at when the file is generated.
  - **Results File**: Select this to download the results file. The file displays the content of the database after the content of the file replaced the content of the database. You can view the results file only after the replace operation is complete.

You can export these files to your local system.

## Creating a Data File for Import

If you are migrating new data into the database, you must prepare the data file using the correct format and syntax before you can import it successfully. You must include all the required fields and understand the dependencies among some of the fields. For detailed information about the guidelines, supported record types, and interdependencies among fields, refer to the [Infoblox CSV Import Reference](#).

## Exporting Data to Files

You can export existing data to a CSV file. The appliance marks all required fields with an asterisk ★ in the exported file. It also adds a `_new_XXXX` field to each required field so you can use this field to update data. You cannot stop an export once you start it.

To export all data to a CSV file:

1. From Grid Manager, navigate to the panel that contains the data you want to export. For example, if you want to export data for all DNS zones, select the **Data Management** tab -> **DNS** tab -> **Zones** tab.
2. In the panel, select **Export data in Infoblox CSV Import format** from the **Export** drop-down menu.
3. In the *Export* dialog box, complete the following:
  - a. **Separator**: Select the separator used in the data file. The default is Comma.
  - b. Click **Export**.

The appliance exports all the fields of the records that are displayed in this panel based on your filter criteria. You can either open the data file or save it to your computer. The appliance uses a default file name depending on the panel from which you perform the export. For example, when you export the data from the **IPAM** tab, the default file name is *Allnetworks.csv*. When you export data from the **DNS** tab, the default file name is *Allzones.csv*. The file contains a header row that includes all the fields of the corresponding record type. You can update this data file, and then re-import the data into the database.

You can also export the displayed fields in a panel.

## Exporting New CSV Jobs

You can use the *Global CSV Export Wizard* to export multiple objects at once. You can export multiple object data types to a single CSV file through Grid Manager. Objects that have the least or no dependencies on other objects are placed at the top of the list and the most dependent objects are displayed at the bottom of the list. The export command will not re-execute after a Grid Manager HA failover.

For standalone appliances, the Grid Manager sends a single CSV export request that contains both Grid and member CSV headers to export Grid and member properties. Similarly, Grid Manager combines GridDns and MemberDns headers to export DNS properties and GridDhcp and MemberDhcp headers to export DHCP properties.

You cannot export the default DNS view if you have not created a custom DNS view. To export the default DNS view, you must either create custom DNS views or use the **CSV Global Export** option.

Note that when you use the **Global CSV Export** option to export either all objects, all DNS objects, or all PTR records, the appliance also exports the PTR system generated records from the auto-created zone "0.0.127.in-addr.arpa". When you use the same file for import operation, the CSV import operation might fail due to the presence of these system generated records. Infoblox recommends that you select the **Skip to next row and continue** option while performing the CSV import operation. When you select this option, CSV import skips these rows with appropriate error messages and then processes other rows in the CSV file.

To export DNS and DHCP data:

1. From the **Data Management** tab, click **CSV Job Manager** from the Toolbar.
2. In the *CSV Job Manager* wizard, click **CSV Export**. Click the New CSV Export job icon and specify the following in the *Global CSV Export Wizard*:
  - **Separator**: Select a separator from the drop-down list: **Comma**, **Semicolon**, **Space**, and **Tab**.
  - **All Objects**: The checkbox is selected by default. When you select this checkbox, all DHCP, global, and DNS objects are selected by default.

- **All DHCP Objects:** Select this checkbox to select all DHCP objects that are listed. To select specific DHCP objects, clear this checkbox and select respective checkboxes next to the DHCP object names.
  - **Global Objects:** Select this checkbox to select all global objects that are listed. To select specific global objects, clear this checkbox and select respective checkboxes next to the global object names.
  - **All DNS Objects:** Select this checkbox to select all DNS objects that are listed. To select specific DNS objects, clear this checkbox and select respective checkboxes next to the DNS object names.
3. Click **Export Data** to start the export process. In the *Start global CSV Export* dialog box, click **Yes** to confirm or click **No** to cancel the process.
  4. Grid Manager displays the CSV export progress and results in the *CSV Global Export progress* dialog box. It displays the following information:
    - **Separator:** The separator used in the CSV file.
    - **Approximate number of objects to be exported:** Indicates the approximate number of objects to be exported.
    - **Number of objects exported:** Indicates the total number of objects that are exported.
    - **Current status:** Indicates the current status of the export process.
  5. Click **Close** to exit.

## Managing CSV Export Jobs

You can view the list of CSV export operations using the **CSV Job Manager**.

1. From the **Data Management** tab, click **CSV Job Manager** from the Toolbar. In the *CSV Job Manager* wizard, click **CSV Export**.
2. Grid Manager displays the following information:
  - **User Name:** The admin user who submitted the CSV export. Only superusers can view this column.
  - **Status:** The current status of the export job. The CSV export process runs in an asynchronous mode similar to the CSV import. The status can be one of the following:
    - **Export pending:** The job is in queue for execution.
    - **Export running:** The job is being executed.
    - **Export stopped:** The job has been stopped. You can select the job and restart the export.
    - **Export completed:** The export is completed without errors. Check the **Message** field for information about the export.
    - **Export failed:** The export is completed, but with errors. Check the **Message** field for information about the error message.
  - **Submitted:** The timestamp when the job was submitted.
  - **Completed:** The timestamp when the job was completed. This field is blank if the job has not been completed yet.
  - **Failed Description:** The appliance displays the error message in the **Failed Description** column if the CSV export fails. You cannot start a failed CSV export task.
  - **File Size:** The size of the generated CSV file.

Completed CSV export jobs are deleted 30 days from the date of submission. You can also perform the following operations:

- Select a completed CSV export task and download the exported file.
- Cancel or stop a pending CSV export task using the **Delete** option.
- Purge all export tasks after an upgrade or restore.
- Start a new global CSV export job.

## Modifying CSV Export Jobs

You can cancel an export job that is in progress or download an exported job. You can also schedule new export jobs. The appliance deletes the CSV jobs that are completed after 30 days from the date of submission.

To update an exported job or delete one:



1. From the **Data Management** tab, click **CSV Job Manager** from the Toolbar.
2. In the *CSV Job Manager* wizard, click **CSV Export**.
3. Select the export job that you want to update, click the Action icon. Click **Cancel** to cancel the export job that is in progress. You can also click the **Cancel job** icon to delete the file. Click **Download file** to download the exported file.
4. Click the New CSV Export job icon to export a new job.
5. Click **Close** to exit.

## Exporting Displayed Data

You can export visible information, such as global search results and the syslog file, in CSV format from panels and pages that support the Export function, and then easily convert the file to PDF and other file formats. You can also export all data in a specific panel.

To export displayed data:

1. From Grid Manager, navigate to the panel that contains the data you want to export. For example, if you want to export data for DNS zones, select the **Data Management** tab -> **DNS** tab -> **Zones** tab.
2. In the panel, select **Export visible data** from the **Export** drop-down menu.
3. In the *Export* dialog box, click **Start**. Grid Manager displays a message about the time required to export data could be long depending on the amount of data.
4. Click **Download** when the export is finished.
5. Depending on your browser and operating system, you may need to do one of the following in the *Opening .csv* dialog box:
  - **Open with:** Select a program with which you want to open the .csv file.
  - **Save to Disk:** Select this if you want to save the .csv file to your local computer.
  - **Do this automatically for files like this from now on:** Select this checkbox if you want Grid Manager to use the same method for future exports. When you select this checkbox, Grid Manager does not display the *Opening .csv* dialog box in the future.
6. Click **OK**.

Depending on the selected option, Grid Manager opens the file using the program you select, or saves the file to your local computer.

## Related topic

[CSV Import Reference](#)

## SSL and TLS Protocols

When you log in to the NIOS appliance, your computer makes an HTTPS (Hypertext Transfer Protocol over Secure Sockets Layer protocol) connection to the NIOS appliance. HTTPS is the secure version of HTTP, the client-server protocol used to send and receive communications throughout the Web. HTTPS uses SSL (Secure Sockets Layer) and/or TLS (Transport Layer Security) protocols to secure the connection between a client and server. SSL/TLS provides server authentication and encryption. The NIOS appliance supports TLS versions 1.0, 1.1, and 1.2. TLS provides cipher suites that are used to negotiate the security settings for the secure connection. Infoblox has provided a few CLI commands so you can enable and disable specific cipher suites. For detailed information about these CLI commands, see [Using the NIOS CLI](#).

Note that enabling or disabling the TLS ciphers will enable or disable the equivalent SSHd cipher. The following table lists the TLS suite name and the corresponding OpenSSL suite name, SSHd cipher name, and SSHd MAC name:

*TLS Cipher Suites*

TLS Suite Name	Open SSL Suite Name	SSHd Cipher	SSHd MAC



TLS_DHE_RSA_WITH_AES_256_CBC_SHA	DHE-RSA-AES256-SHA	aes256-cbc	hmac-sha1, hmac-sha1-etm@openssh.com
TLS_DHE_DSS_WITH_AES_256_CBC_SHA	DHE-DSS-AES256-SHA	aes256-cbc	hmac-sha1, hmac-sha1-etm@openssh.com
TLS_RSA_WITH_AES_256_CBC_SHA	AES256-SHA	aes256-cbc	hmac-sha1, hmac-sha1-etm@openssh.com
TLS_DH_RSA_WITH_3DES_EDE_CBC_SHA	EDH-RSA-DES-CBC3-SHA	3des-cbc	hmac-sha1, hmac-sha1-etm@openssh.com
TLS_DH_DSS_WITH_3DES_EDE_CBC_SHA	EDH-DSS-DES-CBC3-SHA	3des-cbc	hmac-sha1, hmac-sha1-etm@openssh.com
TLS_RSA_WITH_3DES_EDE_CBC_SHA	DES-CBC3-SHA	3des-cbc	hmac-sha1, hmac-sha1-etm@openssh.com
TLS_DHE_DSS_WITH_AES_128_CBC_SHA	DHE-DSS-AES128-SHA	aes128-cbc	hmac-sha1, hmac-sha1-etm@openssh.com
TLS_DHE_DSS_WITH_AES_128_CBC_SHA	DHE-DSS-AES128-SHA	aes128-cbc	hmac-sha1, hmac-sha1- etm@openssh.com
TLS_RSA_WITH_AES_128_CBC_SHA	AES128-SHA	aes128-cbc	hmac-sha1, hmac-sha1-etm@openssh.com
TLS_RSA_WITH_RC4_128_SHA	RC4-SHA	arcfour128	hmac-sha1, hmac-sha1-etm@openssh.com
TLS_DHE_DSS_WITH_AES_256_GCM_SHA384	DHE-DSS-AES256-GCM-SHA384	aes256- gcm@openssh.com	hmac-sha2-512, hmac-sha2-512- etm@openssh.com
TLS_DHE_RSA_WITH_AES_256_GCM_SHA384	DHE-DSS-AES256-GCM-SHA384	aes256- gcm@openssh.com	hmac-sha2-512, hmac-sha2-512- etm@openssh.com
TLS_DHE_RSA_WITH_AES_256_CBC_SHA256	DHE-RSA-AES256-SHA256	aes256-cbc	hmac-sha2-256, hmac-sha2-256- etm@openssh.com
TLS_DHE_DSS_WITH_AES_256_CBC_SHA256	DHE-DSS-AES256-SHA256	aes256-cbc	hmac-sha2-256, hmac-sha2-256- etm@openssh.com
TLS_RSA_WITH_AES_256_GCM_SHA384	AES256-GCM-SHA384	aes256- gcm@openssh.com	hmac-sha2-512, hmac-sha2-512- etm@openssh.com
TLS_DHE_DSS_WITH_AES_128_GCM_SHA256	DHE-DSS-AES128-GCM-SHA256	aes128- gcm@openssh.com	hmac-sha2-256, hmac-sha2-256- etm@openssh.com

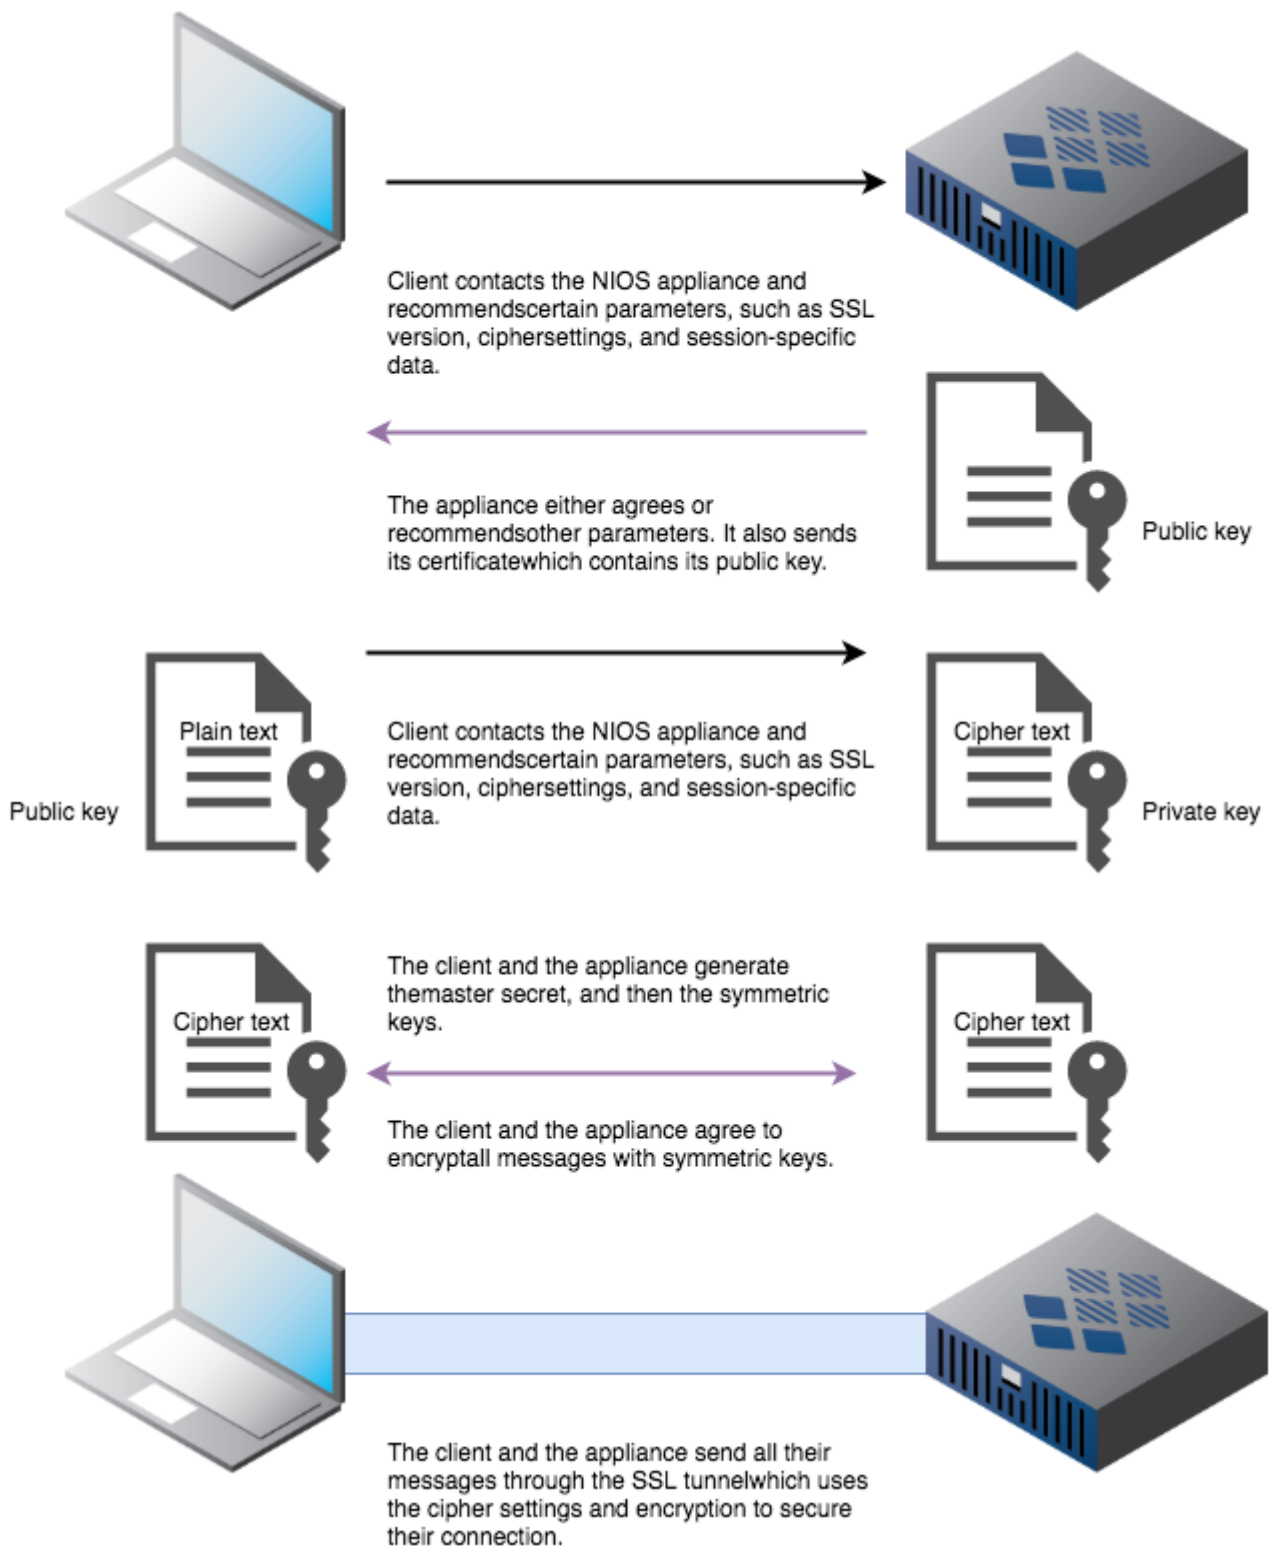
TLS_DHE_RSA_WITH_AES_128_GCM_SHA256	DHE-RSA-AES128-GCM-SHA256	aes128-gcm@openssh.com	hmac-sha2-256, hmac-sha2-256-etm@openssh.com
TLS_DHE_RSA_WITH_AES_128_CBC_SHA256	DHE-RSA-AES128-SHA256	aes128-cbc	hmac-sha2-256, hmac-sha2-256-etm@openssh.com
TLS_DHE_DSS_WITH_AES_128_CBC_SHA256	DHE-DSS-AES128-SHA256	aes128-cbc	hmac-sha2-256, hmac-sha2-256-etm@openssh.com
TLS_RSA_WITH_AES_128_GCM_SHA256	AES128-GCM-SHA256	aes128-gcm@openssh.com	hmac-sha2-256, hmac-sha2-256-etm@openssh.com
TLS_RSA_WITH_AES_128_CBC_SHA256	AES128-SHA256	aes128-cbc	hmac-sha2-256, hmac-sha2-256-etm@openssh.com

When a client first connects to a server, it starts a series of message exchanges, called the SSL/TLS handshake. During this exchange, the server authenticates itself to the client by sending its server certificate. A certificate is an electronic form that verifies the identity and public key of the subject of the certificate. (In SSL/TLS, the subject of the certificate is the server.) Certificates are typically issued and digitally signed by a trusted third party, the Certificate Authority (CA). A certificate contains the following information: the dates it is valid, the issuing CA, the server name, and the public key of the server. For information about certificates, see [Managing Certificates](#).

A server generates two distinct but related keys: a public key and a private key. During the SSL/TLS handshake, the server sends its public key to the client. Once the client validates the certificate, it encrypts a random value with the public key and sends it to the server. The server decrypts the random value with its private key.

The server and the client use the random value to generate the master secret, which they in turn use to generate symmetric keys. The client and server end the handshake when they exchange messages indicating that they are using the symmetric keys to encrypt further communications.

### SSL/TLS Handshake



To avoid possible attacks in which HTTP or HTTPS connections are made to a web server and stay open much longer than they should be, Infoblox provides the `set connection_limit` and `show connection_limit` CLI commands that you can use to mitigate these attacks. In general, these attacks can result in the web server reaching its maximum number of concurrent connections, and thus denying connections from legitimate sources. You can use the

CLI commands to limit the number of concurrent HTTP and HTTPS connections from a given client that corresponds to a particular IP address. For information about the CLI commands and how to use them, see [Using the NIOS CLI](#).

## Managing Certificates

This topic covers the following sections:

- [About Client Certificates](#)
- [About Validate Certificates](#)

## About Client Certificates

This section covers the following:

- [Generating a Client Certificate](#)
- [Viewing Client Certificates](#)
- [Downloading Client Certificates](#)

You can generate client certificates for a Grid Master or a Grid Master candidate, and then send it to another server, such as a Hardware Security Module (HSM).

## Generating a Client Certificate

To generate a client certificate:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab.  
Grid Master Candidate: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox.
2. From the Toolbar, click **Certificates** -> **Client Cert** -> **Generate Client Certificate**, and select either **RSASHA1** or **RSASHA256**.
  - If you are generating a certificate for an HSM group with Thales Luna 4 devices, you must select **RSASHA1**; and if the certificate is for an HSM group with Thales Luna 5 or Luna 6 devices, select **RSASHA256**.

The appliance displays a confirmation dialog after it generates the certificate. If a certificate had been previously generated, the appliance displays a dialog warning that if the previous certificate was registered with a server, then the new certificate must be registered with the server.

## Viewing Client Certificates

To view the client certificates that were generated:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab.  
Grid Master Candidate: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox.
2. From the Toolbar, click **Certificates** -> **Client Cert** -> **View Client Certificate**, and select either **RSASHA1** or **RSASHA256**.

The appliance displays the selected certificate.

## Downloading Client Certificates

To download a client certificate:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab.  
Grid Master Candidate: From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox.
2. From the Toolbar, click **Certificates** -> **Client Cert** -> **Download Client Certificate**, and select either **RSASHA1** or **RSASHA256**.

3. Save the certificate.

## About Validate Certificates

You can view the status results of the certificates uploaded by clicking **Validate Certificates**. The **Validate Certificates** feature validates the following:

- All the certificate files will be validated through OpenSSL command on each file.
- The certificates serial number stored in the database (one.x509\_certificate DB objects), and saves the certificates to temporary files.

To view the validated certificates in Grid Manager:

1. Grid: From the **Grid** tab, click **Certificates** in the Toolbar > **Validate Certificates**. The *Certification Validation Results* dialog displays the following:
  - **Filename/Serial number:** Displays the filename of the certificate files in the folder (For example, omsp\_ca\_cert.pem) or the serial number of the certificates in the database ((For example, 4786438514b4fa8325d750a39ca79.... )
  - **Verification Result:** Displays the certificate result and the certificate status. For example:
    - Valid certificate:  
/storage/tmp/cert1.pem: OK (folder/temporary files/certificate name.pem: OK)
    - Expired certificate:  
/infoblox/security/certs/ocsp\_ca\_cert.pem: C = IN, ST= Karnataka, L = Bengaluru, O = Infoblox, OU= QA, CN = adityagfc33.inca.infoblox.com, emailAddress = adityag@infoblox.com error 10 at O depth lookup:certificate has expired OK  
(folder/temporary files/certificate name.pem: certificate has expired message OK)
    - Invalid certificate  
/infoblox/security/certs/ocsp\_ca\_cert.pem: C = IN, ST = Karnataka, O = Infoblox, OU = QA2, CN = test.inca.infoblox.com error 20 at 0 depth lookup:unable to get local issuer certificate Warning: contains CA certificate(s) without SKI  
(folder/temporary files/certificate name.pem: issuer certificate Warning message stating contains CA certificate(s) without SKI OK)
2. Click **Close** to exit the *Certification Validation Results* dialog.

### Note

- For valid self-signed certificates, Grid Manager does not display any additional information. However, for expired and invalid certificates Grid Manager displays the issuer/subject information in the *Certification Validation Results* screen.
- Warning messages are displayed for expired certificates and for certificates with no SKI (Subject Key Identifier) in the concatenated certificate file.

## Configuring Approval Workflows

Approval workflows support routing certain core network service tasks submitted by an admin group to another for approval. You can add an admin group to an approval workflow and define the group as a submitter or approver group. Note that only superusers can create approval workflows. For information about how to set up admin groups, see [About Admin Groups](#).

In an approval workflow, you can add a submitter group and an approver admin group that you have previously defined. You can also define when and to whom email notifications are sent, and configure options such as whether submitters or approvers must enter a comment or a ticket number when they submit tasks for approval. Approval workflows are useful when you want to control tasks that require reviews. For example, if you have a group of help desk users who can add, modify, and delete hosts and you want members of an operation group to review these tasks, you can define the help desk users as submitters, and then set up members of the operation group as approvers. You can then add the submitter and approver groups to an approval workflow and configure notifications options and other configurations, such as allowing the approvers to reschedule the submitted tasks.

Not all core network service tasks can be routed for approval. You can configure approval tasks associated with certain objects. For a list of supported objects, see [Scheduling Tasks](#).



### Note

When an admin group is defined as a submitter group, there are certain operations the submitters cannot perform even though they may have the permissions to do so. For information about such operations, see [Unsupported Operations for Submitters](#) below.

To create an approval workflow, complete the following:

1. If you have not already done so, set up admin groups that you can configure as submitter groups and approver groups in an approval workflow, as described in [About Admin Groups](#).
2. Create an approval workflow and configure email notifications and other options, as described in [Creating Approval Workflows](#) below.

You can do the following after you have created approval workflows:

- View a list of approval workflows, as described in [Viewing Approval Workflows](#) below.
- Modify approval workflows, as described in [Modifying Approval Workflows](#) below.
- Delete approval workflows, as described in [Deleting Approval Workflows](#) below.
- View a list of approval tasks, as described in [Viewing Approval Tasks](#) below.
- View approval notifications, as described in [Viewing Workflow Notifications](#) below.

## Supported Tasks for Different Admin Groups

Depending on your admin permissions, you may or may not be able to perform certain tasks that are subject to approvals. Supported Tasks for Different Admin Groups lists specific tasks and indicates which admin group can perform the tasks.

### *Supported Tasks for Admin Groups*

Admin groups that can perform the task			
Tasks related to approval workflows	Submitters	Approvers	Superusers
Change the schedule of a task when it is pending approval	Yes	No	Yes
Change the schedule of a task after it has been approved	Yes (Task is re-submitted for approval)	Yes	Yes
Execute the task now when it is pending approval	No	No	No
Delete a task after it has been approved but pending execution	No	No	Yes
Delete a task after it failed or has been executed	No	No	Yes
Delete tasks by selecting the <b>Select all objects in this dataset</b> option	Yes	Yes	Yes

Admin groups that can perform the task			
Tasks related to approval workflows	Submitters	Approvers	Superusers
<hr/> <p><b>Note:</b> Not all tasks are deleted, depending on the task status and the admin who performs the deletion</p> <hr/>			

## Creating Approval Workflows

Before you create an approval workflow, ensure that you have admin groups that you can define as submitters and approvers. Note that a submitter group can be added to only one approval workflow, and approver groups can be added to multiple workflows. An approver can choose to approve a task and either keep or change the date and time when the task is executed. For information about scheduling and rescheduling tasks, see [Scheduling Tasks](#). An approver can also reject a submitted task.

All submitted tasks are executed based on submitter permissions. When an admin submits a task, the appliance logs the task in the audit log and associates it with a task ID. You can view your tasks in **Task Manager**, as described in [Viewing Tasks](#). Depending on your configuration, you can control when and to whom email notifications are sent. For example, you can configure the appliance to send notifications to only the approver each time when a task requires approval, or send notifications to both the submitter and approver group each time when a task is disapproved.

To create an approval workflow

- From the **Administration** tab, select the **Workflow** tab -> **Approval Workflows** tab, and then click the Add icon
- In the *Add Approval Workflow* wizard, complete the following:
  - Submitter Group:** From the drop-down list, select the admin group whose submitted tasks require approvals. Note that performing CSV imports do not require approvals. If there is a warning that the submitter group has CSV import permission, you may want to remove the permission.
  - Approver Group:** From the drop-down list, select the group that can approve tasks submitted by admins of the submitter group. If the approver group you select does not have the permission to schedule tasks, the approvers cannot reschedule the execution dates and times of the tasks when they approve them.
  - Ticket Number:** From the drop-down list, select one of the following to determine whether a ticket number is required when a submitter submits a task for approval.
    - Required:** The submitter must enter a ticket number when submitting a task.
    - Optional:** The submitter can choose to enter a ticket number or not when submitting a task.
    - Not Used:** The **Ticket Number** field does not appear when the submitter creates a task.
  - Submitter Comment:** From the drop-down list, select whether the submitter must enter a comment or not when submitting a task for approval. You can select **Required**, **Optional**, or **Not Used**.
  - Approver Comment:** From the drop-down list, select whether the approver must enter a comment or not when approving a task. You can select **Required**, **Optional**, or **Not Used**.
- Click **Next** and complete the following to specify notification options for the workflow:
  - Approver Notification Address(es):** Select one of the following to specify to which approver email addresses the appliance sends workflow notifications. The default is **Group Email Address(es)**.
    - Group Email Address(es):** Select this if you want the appliance to send notifications to the list of email addresses configured for the admin group. For information about how to configure this list, see [About Admin Groups](#).

- **User Email Address(es)**: Select this if you want the appliance to send notifications to individual email addresses of the admin group.
- **Notifications sent on**: Select the operations that can trigger email notifications. When you select an operation, the appliance sends a notification each time that operation occurs. By default, all operations are selected.
  - **Approval Required**: The appliance sends an email notification each time an approval is required.
  - **Task Approved**: The appliance sends an email notification each time a task is approved.
  - **Task Rejected**: The appliance sends an email notification each time a task is rejected.
  - **Task Succeeded**: The appliance sends an email notification each time a task is completed successfully.
  - **Task Failed**: The appliance sends an email notification each time the execution of a task fails.
  - **Task Rescheduled**: The appliance sends an email notification each time a task is being rescheduled.
- **Notifications sent to**: For each operation, select whether the **Approver**, **Submitter**, or **Both** are notified when the operation occurs. The default value is **Both** for all operations. For information about email notifications, see Viewing Approval Tasks below.

4. Optionally, click **Next** to add extensible attributes to the approval workflow. For information, see [About Extensible Attributes](#).

5. Save the configuration.

## Viewing Approval Workflows

Grid Manager lists all approval workflows in the **Approval Workflows** tab. Only superusers can view approval workflows defined for the Grid. Limited-access users cannot view approval workflows.

To view approval workflows:

1. From the **Administration** tab, select the **Workflow** tab -> **Approval Workflows** tab.
2. Grid Manager displays the following for each approval workflow:
  - **Submitter Group**: The name of the admin group whose tasks require approvals.
  - **Approver Group**: The name of the admin group that can approve tasks submitted by members of the submitter group.
  - **Ticket Number**: Displays whether the submitter is required to enter a ticket number when submitting tasks that require approvals. Possible values are **Not Used**, **Optional**, and **Required**.
  - **Submitter Comment**: Displays whether the submitter is required to enter a comment when submitting tasks that require approvals. Possible values are **Not Used**, **Optional**, and **Required**.
  - **Approver Comment**: Displays whether the approver is required to enter a comment when approving tasks. Possible values are **Not Used**, **Optional**, and **Required**.
  - **Site**: Values that were entered for this predefined extensible attribute.

You can do the following in this tab:

- Modify some of the data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes. Note that some fields are read-only.
- Sort the data in ascending or descending order by column.
- Select an approval workflow and click the Edit icon to modify data, or click the Delete icon to delete it.
- Use filters and the **GoTo** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information about using quick filters, see [Finding and Restoring Data](#).
- Print and export the data in this tab.

## Modifying Approval Workflows

You can modify information in an approval workflow, except for the submitter group. To modify approval workflow configuration:

1. From the **Administration** tab, select the **Workflow** tab -> **Approval Workflows** tab.
2. Select an approval workflow and click the Edit icon.
3. Grid Manager provides the following tabs from which you can modify information:



- **General** tab: You can modify the approver group and decide whether the ticket number, submitter comment, and approver comment are required, but you cannot change the submitter group.
  - **Approval Notifications** tab: You can modify when and to whom email notifications are sent.
  - **Extensible Attributes** tab: You can add or modify values of extensible attributes. For information, see [About Extensible Attributes](#).
4. Save the configuration.

## Deleting Approval Workflows

You can delete an approval workflow any time after you have created it. Note that when you delete a workflow that has associated tasks that are pending approvals, the tasks will be rejected after you delete the workflow.

To delete an approval workflow:

1. From the **Administration** tab, select the **Workflow** tab -> **Approval Workflows** tab.
2. Select an approval workflow and click the Delete icon.
3. Click **Yes** in the *Delete Confirmation* dialog.

## Viewing Approval Tasks

If you belong to an approver admin group, you can view, approve, or reject tasks that are pending your approval in the **Task Manager** tab. For information, see [Viewing Tasks](#). Submitter's can view all pending and completed tasks they have submitted.

## Viewing Workflow Notifications

When a submitter and approver receives an email notification about their tasks, the appliance lists the approval status and workflow related information such as task ID, submitter name, execution time, object type and action in the email notification.

Following is a sample email notification:

Notification:

=====

Message: Task 32 submitted by subm has been approved The following task has been approved:

Task details

Task ID: 32 Submitter: subm Approver: jdoe

Submit time: 2012-10-09 05:55:01 (UTC) Coordinated Universal Time

Execution time: N/A

Object type: NS Record Action: Add

Affected object: corp1.com Ticket number: MKTG245

Submitter comment: Create an NS record. Approver comment: Approved.

Click here to go to the task management tab - <https://192.168.1.2/ui/?contextId=taskmanager>



### Note

When you can click the hyperlink displayed in the notification, you can log in to **Grid Manager** and access the **Task Manager** tab in a separate browser tab or window.

## Unsupported Operations for Submitters

When admins are part of a submitter group in an approval workflow, there are certain operations they cannot perform even though they may have the permissions to do so. Following is the list of operations that submitters cannot perform:

- Reclaim IPv4 or IPv6 addresses
- Expand networks
- Resize networks
- Split networks
- Sign (DNSSEC) zones
- Unsign DNSSEC signed zones
- Import DS to DNSSEC signed zones
- Perform KSK rollovers on a DNSSEC signed zones
- Copy records from one DNS zone to another
- Clear all discovered data
- Clear discovered timestamps
- Clear unmanaged addresses
- Resolve discovery conflicts
- Update extensible attributes on multiple objects at a the same time
- Delete or modify several objects at a time (using the "Select all objects in this dataset" option from Grid Manager)
- Order DHCP Ranges inside a network (feature is available only when used with Sophos)
- Configure member DHCP Captive Portal through the wizard
- Restore objects from the Recycle Bin
- Delete non-native NIOS DNS resource records. These objects can only be synchronized from a Microsoft DNS server
- Copy rules from one Response Policy Zone to another
- Order Response Policy Zones

## Administering NIOS

This section provides information about configuring admin groups, roles, and accounts, and defining the appropriate permissions. It explains how to configure and manage a Grid or an independent appliance, and set operational parameters. It also describes the file distribution services (TFTP, FTP, and HTTP) and the bloxTools environment. It includes the following topics:

- [Appliance Administration](#)
- [IP Address Management](#)
- [Configuring Super Hosts](#)
- [DNS](#)
- [DHCP](#)
- [Configuring Microsoft Windows Servers](#)
- [Monitoring](#)
- [Infoblox Reporting and Analytics](#)
- [Infoblox Infrastructure Security](#)
- [Infoblox Subscriber Services](#)
- [VLAN Management](#)

## Appliance Administration

This section provides information about configuring admin groups, roles, and accounts, and defining the appropriate permissions. It explains how to configure and manage a Grid or an independent appliance, and set operational parameters. It also describes the file distribution services (TFTP, FTP, and HTTP) and the bloxTools environment. It includes the following chapters:

- [Managing Administrators](#)
- [Deploying a Grid](#)
- [Deploying Independent Appliances](#)
- [Deploying Cloud Network Automation](#)
- [Managing Appliance Operations](#)
- [File Distribution Services](#)
- [bloxTools Environment](#)
- [RIR Registration Updates](#)

## Managing Administrators

This section describes the various tasks associated with setting up admin groups, admin roles, admin accounts, and permissions. It contains the following sections:

- [About Admin Groups](#)
- [About Admin Roles](#)
- [Managing Admin Groups and Admin Roles](#)
- [About Administrative Permissions](#)
- [Authenticating Administrators](#)
- [Authenticating Admins Using Two-Factor Authentication](#)
- [About Remote Admins](#)
- [Authenticating Admins Using RADIUS](#)
- [Creating Local Admins](#)
- [Authenticating Admin Accounts Using TACACS+](#)
- [Authenticating Admins Using LDAP](#)
- [About Admin Accounts](#)
- [Defining the Authentication Policy](#)
- [Authenticating Admins Using Active Directory](#)

- [Changing Password Length Requirements](#)
- [Notifying Administrators](#)
- [Administrative Permissions for Common Tasks](#)
- [Administrative Permission for the Grid](#)
- [Administrative Permissions for IPAM Resources](#)
- [Administrative Permissions for DNS Resources](#)
- [Administrative Permissions for DNS Resources with Associated IP addresses in Networks and Ranges](#)
- [Administrative Permissions for DHCP Resources](#)
- [Administrative Permissions for File Distribution Services](#)
- [Administrative Permissions for Dashboard Tasks](#)
- [Administrative Permissions for Certificate Authentication Services and CA Certificates](#)
- [Administrative Permissions for Object Change Tracking](#)
- [Administrative Permissions for Load Balancers](#)
- [Authenticating Admins Using SAML](#)
- [Administrative Permissions for Named ACLs](#)
- [Administrative Permissions for DNS Threat Protection](#)
- [Administrative Permissions for Cloud Objects](#)
- [Administrative Permissions for Reporting](#)
- [Administrative Permissions for VLAN Management](#)
- [Administrative Permissions for SAML](#)

## About Admin Groups

All administrators must belong to an admin group. The permissions and properties that you set for a group apply to all administrators assigned to that group. You can assign a dashboard template to an admin group. A dashboard template specifies the tasks an admin group can access through the **Tasks Dashboard** tab when they log in to Grid Manager. For information about dashboard templates, see [Configuring Dashboard Templates](#). You can also restrict certain user groups to manage specific tasks in the **Tasks Dashboard** tab only. These users cannot manage other core network services through Grid Manager. For information about how to apply this restriction, see [Limited-Access Admin Groups](#) below.

To define admins who can perform specific core network service tasks, you can set up admin groups and assign them permissions for those tasks. To control when and whether certain tasks should be performed, you can add an admin group to an approval workflow and define the admins as submitters or approvers. A submitter is an admin whose tasks require approvals before execution, and an approver is an admin who can approve the submitted tasks. When you add submitter and approver groups to an approval workflow, you have control over who can perform which mission critical tasks and whether and when the tasks should be executed. For more information about how to create and configure approval workflows, see [Configuring Approval Workflows](#).

There are three types of admin groups:

- **Superuser** – Superuser admin groups provide their members with unlimited access and control of all the operations that a NIOS appliance performs. There is a default superuser admin group, called **admin-group**, with one superuser administrator, **admin**. You can add users to this default admin group and create additional admin groups with superuser privileges. Superusers can access the appliance through its console, GUI, and API. In addition, only superusers can create admin groups.
- **Limited-Access** – Limited-access admin groups provide their members with read-only or read/write access to specific resources. These admin groups can access the appliance through the GUI, API, or CLI. They cannot access the appliance through the console.
- **Default** – When upgrading from previous NIOS releases, the appliance converts the ALL USERS group to the Default Group when the ALL USERS Group contains admin accounts. The appliance does not create the Default Group if there is no permission in the ALL USERS group. The permissions associated with the ALL USERS group are moved to a newly created role called Default Role. Supported in previous NIOS releases, the ALL USERS group was a default group in which you defined global permissions for all limited-access users. This group implicitly included all limited-access users configured on the appliance.

All limited-access admin groups require either read-only or read/write permission to access certain resources, such as Grid members, and DNS and DHCP resources, to perform certain tasks. Therefore, when you create an admin group, you must specify which resources the group is authorized to access and their level of access.

Only superusers can create admin groups and define their administrative permissions. There are two ways to define the permissions of an admin group. You can create an admin group and assign permissions directly to the group, or you can create roles that contain permissions and assign the roles to an admin group.

You must create admin groups and assign them access to the cloud API and applicable permissions so they have authority over delegated objects. When you assign permissions for objects that have not been delegated, these admin groups or admin users assume applicable permissions to these un-delegated objects. For example, you can create an admin group that can access a specific set of networks while another can access another set of networks. Note that you cannot create a new admin group using the same name. For information about Cloud Network Automation, see [Deploying Cloud Network Automation](#).

Complete the following tasks to assign permissions directly to an admin group:

1. Create an admin group, as described in [Creating Limited-Access Admin Groups](#) below.
2. Assign permissions to the admin group, as described in [About Administrative Permissions](#). Complete these tasks to assign admin roles to an admin group:
3. Create an admin role, as described in [About Admin Roles](#).
4. Define permissions for the newly created admin role, as described in [Creating Admin Roles](#), see [About Admin Roles](#).
5. Create an admin group and assign the role to the group, as described in [Creating Limited-Access Admin Groups](#) below.

After you have created admin groups and defined their administrative permissions, you can assign administrators to the group.

- For local admins, see [Creating Local Admins](#).
- For remote admins, see [About Remote Admins](#).

## Creating Superuser Admin Groups

Superusers have unlimited access to the NIOS appliance. They can perform all operations that the appliance supports. There are some operations, such as creating admin groups and roles, that only superusers can perform.

Note that there must always be one superuser admin account, called "admin", stored in the local database to ensure that at least one administrator can log in to the appliance in case the NIOS appliance loses connectivity to the remote admin databases such as RADIUS servers, AD domain controllers, TACACS+ servers, LDAP servers, or OCSP responders.

NIOS comes with a default superuser admin group (**admin-group**). It also automatically creates a new admin group, **fireeye-group**, when you add the first FireEye RPZ (Response Policy Zone). Infoblox recommends that you do not add another admin group with the same name as the default or FireEye admin group. Note that the FireEye admin group is read-only and you cannot assign permissions to it. For more information about FireEye RPZs, see [About FireEye Integrated RPZs](#).

When you install valid licenses and configure your Grid for Cloud Network Automation, NIOS enables the **cloud-api-only** admin group. You can assign admin users to this group so they are authorized to send cloud API requests to the Cloud Platform Appliances. Note that you cannot delete this admin group or create a new admin group using the same name. For information about Cloud Network Automation, see [Deploying Cloud Network Automation](#).

You can create additional superuser admin groups, as follows:

1. From the **Administration** tab, select the **Administrators** tab -> **Groups** tab, and then click the Add icon.
2. In the *Add Admin Group* wizard, complete the following:
  - **Name:** Enter a name for the admin group.
  - **Comment:** Enter useful information about the group, such as location or department. For **fireeye-group**, NIOS displays **Group used to receive FireEye alerts** in this field.

- **Disable:** Select this to retain an inactivated profile for this admin group in the configuration. For example, you may want to define a profile for recently hired administrators who have not yet started work. Then when they do start, you simply need to clear this checkbox to activate the profile.
3. Click **Next** and complete the following:
    - Superusers:** Select this to grant the admin accounts that you assign to this group full authority to view and configure all types of data and perform all tasks.
    - **Allowed Interfaces:** Superusers admin groups are automatically granted access to the Infoblox GUI (Grid Manager), API, and CLI. You can specify which API the superusers group can access. Note that you must have the Cloud Network Automation or Cloud Platform license installed to configure access to the cloud API.
      - GUI:** This is selected by default. The superusers admin group automatically has full access to Grid Manager.
      - API:** This is selected by default. Note that the following options are displayed only if a cloud license is installed in the Grid.
        - CLI:** This is selected by default. The superusers admin group automatically has full access to the NIOS CLI.
    - **API (WAPI/PAPI only):** The superusers admin group has full access to the RESTful API and the Infoblox API by default.
    - **Cloud API:** Select this to allow the superusers admin group to use the cloud API. This option is available only if a cloud license is installed in the Grid. Select one of the following:
      - **Cloud API only (no PAPI):** Select this to allow the admin group to use WAPI (RESTful API) to send cloud API requests. Note that the Cloud API uses WAPI exclusively. The group has no access to the Infoblox API.
      - **Cloud API and PAPI (No WAPI):** Select this to allow the admin group to send API requests and have access to the Infoblox API. However, the group cannot use WAPI to send cloud API calls.



#### Note

When you assign cloud API access to an admin group, the group assumes full authority over all delegated objects. You must however specifically assign object permissions to the admin group for the group to gain authority over non-delegated objects. For information about how to assign object permissions, see [About Administrative Permissions](#).

4. Click **Next** and complete the following to define the dashboard template:
  - **Dashboard Template:** From the drop-down list, select the dashboard template you want to assign to this superuser group. When you apply a dashboard template to an admin group, the template applies to all users in the group. The default is **None**, which means that users in this group can access all licensed tasks in the **Tasks Dashboard** tab if they have the correct permissions to the task-related objects. Note that if you want to delete a template, you must first unassign the template from an admin group, or select **None**, before you can delete it. For more information about dashboard templates, see [About Dashboards](#).
5. Click **Next** to add admin email addresses if you want the appliance to send approval workflow notifications to a list of email addresses for the admin group. Complete the following in the Email Address table:

Click the Add icon and Grid Manager adds a row to the table. Enter the email address of the admin who should receive workflow notifications. You can click the Add icon again to add more email addresses. You can also select an email address and click the Delete icon to delete it. To modify an email address, click the **Email Address** column and modify the existing address.



#### Note

When you configure an approval workflow and select **Group Email Address(es)** as the approver notification addresses, the appliance sends workflow notifications to all email addresses you have added to this table. For information, see [Configuring Approval Workflows](#).

6. Optionally, click **Next** to add extensible attributes to the admin group. For information, see [About Extensible Attributes](#).

7. Save the configuration and click **Restart** if it appears at the top of the screen. You can do one of the following after you create a superuser admin group:

- Add local admins to the superuser group. For information, see [Creating Local Admins](#).
- Assign the superuser group to remote admins. For information, see [About Remote Admins](#).

## Creating Limited-Access Admin Groups

When you create a limited-access admin group, you can assign roles to it. The group then inherits the permissions of its assigned roles. In addition, you can assign permissions directly to the group. Only superusers can create admin groups.

To create a limited-access admin group:

1. From the **Administration** tab, select the **Administrators** tab -> **Groups** tab, and then click the Add icon.
2. In the *Add Admin Group* wizard, complete the following:
  - **Name:** Enter a name for the admin group.
  - **Comment:** Enter useful information about the group, such as location or department.
  - **Disable:** Select this to retain an inactivated profile for this admin group in the configuration. For example, you may want to define a profile for recently hired administrators who have not yet started work. Then when they do start, you simply need to clear this checkbox to activate the profile.
3. Click **Next** and complete the following:
  - **Superusers:** Clear this checkbox to create a limited-access admin group.
  - **Roles:** Optionally, click the Add icon to add an admin role to the admin group. In the *Role Selector* dialog box, select the roles you want to assign to the admin group, and then click the Select icon. Use Shift+click and Ctrl+click to select multiple admin roles. You can assign up to 21 roles to an admin group. The appliance displays the selected roles in the list box. When an admin group is assigned multiple roles, the appliance applies the permissions to the group in the order the roles are listed. Therefore if there are overlapped permissions among the roles, the appliance uses the permission from the role that is listed first and ignores the others. You can reorder the list by selecting a role and clicking the arrow keys to move the role up and down the list. To delete a role, select it and click the Delete icon.
  - **Allowed Interfaces:** Specify whether the admin group can use the Infoblox GUI (Grid Manager) and the API (application programming interface) to configure the appliance. Note that you must have the Cloud Network Automation or Cloud Platform license installed to configure access to the cloud API.

**GUI:** Select this to allow the admin group to use the Infoblox GUI, Grid Manager.

**CLI:** Select this to allow the admin group access to the Infoblox CLI. You can select all the commands that you want the group to execute by selecting the command group from the drop-down list. You can then select individual commands from the command group that the admin group can execute. For example, if you want to grant access to the admin group to run all commands related to the Grid command group, select **Grid** from the drop-down list and select all the commands. You can also select individual commands from the Grid command group that you want the admin group to execute.

**API:** Select this to allow the admin group access to the Infoblox API. The following options are available only if a Cloud Network Automation or Cloud Platform license is active in the Grid. You must first select this option to enable the following options.

- **API (WAPI/PAPI only):** Select this to allow the admin group to use only the RESTful API and Infoblox API.
- **Cloud API:** Select this to allow the admin group to use the cloud API. This option is available only if a Cloud Network Automation or Cloud Platform license is installed in the Grid. Select one of the following:
  - **Cloud API only (No PAPI):** Select this to allow the admin group to use WAPI (RESTful API) to send cloud API requests. Note that the Cloud API uses WAPI exclusively. The group has no access to the Infoblox API.
  - **Cloud API and PAPI (No WAPI):** Select this to allow the admin group to send API requests and have access to the Infoblox API. However, the group cannot use RESTful API to send cloud API calls.





#### Notes

- When you assign cloud API access to an admin group, the group assumes full authority over all delegated objects. You must however specifically assign object permissions to the admin group for the group to gain authority over non-delegated objects. For information about how to assign object permissions, see [About Administrative Permissions](#).
- The GUI permission that you assign to the admin group is independent of the CLI permission that you assign. That is, you have to assign each of these permissions separately to non-super users. You can track actions and commands of non-super-users in the audit.log file.
- If you enable CLI commands for reporting users, they will not be able to login to the CLI unless they log in to the **Reporting** tab in Grid Manager.
- SAML-only users will not be able to run CLI commands, because such users are created dynamically and hence do not have the password. However, users belonging to the `saml_local` group can run the `set` series of commands.
- Cloud users will not be able to run CLI commands because they are delegated users.

4. Click **Next** and complete the following to define the dashboard template:

- **DashboardTemplate:** From the drop-down list, select the dashboard template you want to assign to this superuser group. When you assign a dashboard template to an admin group, the template applies to all users in the group. The default is **None**, which means that users in this group can perform all licensed tasks in the **TasksDashboard** tab if they have the correct permissions to the task-related objects. Note that if you want to delete a template, you must first unassign the template from an admin group, or select **None**, before you can delete it. For more information about dashboard templates, see [About Dashboards](#).
- **Display Task flow Dashboards Only:** Select this checkbox if you want to restrict this admin group to access only the Tasks Dashboard in Grid Manager. Note that when you select this checkbox, users in this admin group have access to the tasks you specified in the selected dashboard template, if applicable. They cannot perform any other tasks or manage any core network services in Grid Manager the next time they log in to the system.

5. Click **Next** to add admin email addresses if you want the appliance to send approval workflow notifications to a list of email addresses for the admin group. Complete the following in the Email Address table:

Click the Add icon and Grid Manager adds a row to the table. Enter the email address of the admin who should receive workflow notifications. You can click the Add icon again to add more email addresses. You can also select an email address and click the Delete icon to delete it. To modify an email address, click the **Email Address** column and modify the existing address.



#### Note

When you configure an approval workflow and select **Group Email Address(es)** as the approver notification addresses, the appliance sends workflow notifications to all email addresses you have added to this table. For information, see [Creating Approval Workflows](#).

6. Optionally, click **Next** to add or delete extensible attributes for this admin group. For information, see [About Extensible Attributes](#).

7. Save the configuration and click **Restart** if it appears at the top of the screen.

## About Admin Roles

An admin role is a group of permissions that you can apply to one or more admin groups. Roles allow you to quickly and easily apply a suite of permissions to an admin group. You can define roles once and apply them to multiple admin groups. The appliance contains the following system-defined admin roles:



- **DHCP Admin:** Provides read/write access to all network views, all DHCP MAC filters, all Grid members, and all Microsoft servers that are managed by the Grid. It also provides read-only access to all DHCP templates and DHCP lease history.
- **DNS Admin:** Provides read/write access to all Grid members, all Microsoft servers that are managed by the Grid, all shared record groups, and all DNS views.
- **DTC Admin:** Provides read/write access to all DTC objects, such as LBDNs, LBDN records, pools, servers, monitors, certificates, GeolIP, and topologies.
- **File Distribution Admin:** Provides read/write access to Grid file distribution properties.
- **Grid Admin:** Provides read/write access to all DNS views, all shared record groups, all members, all Microsoft servers that are managed by the Grid, all network views, all DHCP MAC filters, all DHCP templates, DHCP lease history, Grid File distribution properties, network discovery, task scheduling, and all Dashboard tasks.
- **Load Balancer Admin:** Provides read/write access to all load balancer resources.
- **PKI Admin:** Provides read/write access to all HSM groups, all certificate authentication services, and all CA certificates.
- **DHCP Fingerprint:** Provides read/write access to all DHCP fingerprint related objects.
- **Super Host Admin:** Provides read/write access to all super host objects.
- **VLAN Admin:** Provides read/write access to all VLAN views, ranges, and objects.

You can assign these system-defined roles to admin groups and create additional roles based on the job functions in your organization. If you are creating a role that has similar permissions to an existing role, you can copy the role and then make the necessary modifications to the new role. Thus you do not have to create each new role from scratch.

You can assign up to 21 roles to an admin group, and you can assign a role to more than one admin group. When you make a change to a role, the appliance automatically applies the change to that role in all admin groups to which the role is assigned.

## Creating Admin Roles

There are two ways to create an admin role. You can create a new role and define its permissions, or you can copy an existing role and redefine the configuration for the new role.

To create a new role from scratch:

1. From the **Administration** tab, select the **Administrators** tab -> **Roles** tab, and then click the Add icon.
2. In the *Add Role* wizard, complete the following:
  - **Name:** Enter a name for the role.
  - **Comment:** Enter useful information about the role. For example, if you are creating a role for IT personnel, you can put the information here.
  - **Disable:** Select this to retain an inactivated profile for this admin role in the configuration.
3. Optionally, click **Next** to add extensible attributes to this role. For information, see [About Extensible attributes](#).
4. Click **Next** and select one of the following:
  - **Save & Add Permissions:** Save the entry and add permissions to the role. Grid Manager displays the **Permissions** tab with the newly created role selected. You can then add permissions to this role. For information, see [About Administrative Permissions](#).
  - **Save & Close:** Save the entry and close the wizard.
  - **Save & Edit:** Save the entry and continue to edit.
  - **Save & New:** Save the entry and open a new wizard.

To copy an existing role:

1. From the **Administration** tab, select the **Administrators** tab -> **Roles** tab -> *admin\_role* checkbox, and then click **Clone** from the Toolbar.
2. The *Copy Role* editor provides the following tabs from which you can modify data for the new role:
  - **General:** Enter the name and information about the new role. You can also disable the role in this tab.
  - **Admin Groups:** Displays a list of admin groups that are currently using this role. You cannot modify the list.
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with the admin role. You can also modify the values of the extensible attributes. For information, see [About Extensible Attributes](#).
3. Save the configuration and click **Restart** if it appears at the top of the screen. The appliance displays the new role in the **Roles** tab.

After you create roles, you can do the following:

- Define their permissions. For information and guidelines on defining permissions, see [About Administrative Permissions](#).
- Assign roles to admin groups, as described in creating limited-access admin groups, see [About Admin Groups](#).

## Managing Admin Groups and Admin Roles

After you create an admin group or an admin role, you can view, modify, and delete it.

### Modifying Admin Groups and Roles

To modify an admin group:

1. From the **Administration** tab, select the **Administrators** tab -> **Groups** tab -> *admin\_group* checkbox, and then click the Edit icon.
2. The *Admin Group* editor provides the following tabs from which you can modify data:
  - **General**: You can modify the following data.
    - **Name**: Modify the name of the admin group.
    - **Comment**: Enter useful information about the group, such as location or department.
    - **Disable**: Select this to retain an inactivated profile for this admin group in the configuration. For example, you may want to define a profile for recently hired administrators who have not yet started work. Then when they do start, you simply need to clear this checkbox to activate the profile.
    - **Allow Access from**: To control access to the GUI and API, select one of the following. You can restrict access using a named ACL or define individual ACEs. For information about named ACLs, see [Configuring Access Control](#).  
Note this group-based authentication is applicable for Grid-wide settings only. NIOS authenticates user credentials only after it authenticates the Grid-wide settings.
    - **Any**: Select this to allow any clients to access the GUI and API. This is selected by default.
    - **Named ACL**: Select this and click **Select Named ACL** to select a named ACL that contains only IPv4 and IPv6 addresses and networks. When you select this, the appliance allows GUI and API access for all ACEs in the named ACL. You can click **Clear** to remove the selected named ACL.
    - **Set of ACEs**: Select this to configure individual access control entries (ACEs). You can define ACEs for selected admin groups from which users can log in to the application. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding.
    - **IPv4 Address and IPv6 Address**: Select this to add an IPv4 address or an IPv6 address. The **Type** column displays either IPv4 address or IPv6 address based on the item you select from the drop-down list. Click the **Value** field and enter the IP address. The appliance allows this client to access the GUI and API and restricts others.
    - **IPv4 Network and IPv6 Network**: Select this to add an IPv4 network or IPv6 network. The **Type** column displays either IPv4 address or IPv6 address based on the item you select from the drop-down list. Click the **Value** field and enter the network. The appliance allows this network to access the GUI and API and restricts others.
  - After you have added access control entries, you can do the following:
    - Select the ACEs that you want to consolidate and put into a new named ACL. Click the **Create new named ACL** icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
    - Reorder the list of ACEs using the up and down arrows next to the table. Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
  - **Roles**: Modify the data as described in [Creating Limited-Access Admin Groups](#), see [About Admin Groups](#).
  - **Extensible Attributes**: Add and delete extensible attributes that are associated with the admin group. You can also modify the values of the extensible attributes. For information, see [About Extensible Attributes](#).
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring Account Lockout for Admin Groups

You can enable password security such that if a group user tries to log in to Grid Manager by using an incorrect password, NIOS locks the account for a configured time period after the configured number of failed login attempts. Only superusers can enable and configure this feature.

To configure account lockout for admin groups:

1. From the **Administration** tab, select the **Administrators** tab -> **Groups** tab -> *admin\_group* checkbox, and then click the Edit icon
2. In the *Admin Group* editor, select the **Security** tab -> **Basic** tab.
3. NIOS automatically populates some field values based on account lockout configurations for the Grid. For more information, see [Managing Security Operations](#). Click the **Override** button to modify the following data:
  - **Enable Account Lockout:** Select the checkbox to enable account lockout for the group user. This option is disabled by default.
  - **Maximum number of attempts:** Enter the maximum number of invalid login attempts to Grid Manager after which NIOS locks the account. You can specify a value from **1** to **99**. The default value is **5**.
  - **Lockout duration:** Enter the time duration in minutes for which the account must be locked. You can specify a value from **1** to **1440**. The default value is **5**.
  - **Never Unlock:** Select the checkbox to permanently lock a group user account, which is already locked. Only a superuser can clear the checkbox to unlock the account. This option is not applicable to superuser accounts because you cannot permanently lock a superuser account. This option is disabled by default.



### Note

NIOS displays an error on **Save & Close**, if the **Never Unlock** option is enabled for superusers.

## Deleting Admin Groups and Roles

You can remove any default or custom admin group as long as it is not your own admin group or the last admin group. You can also delete any default or custom admin role. The appliance puts the deleted roles in the Recycle Bin, if enabled.



### Note

You cannot delete the **cloud-api-only** and **splunk-reporting-group** admin groups. These admin groups are automatically created when you configure your Grid for Cloud Network Automation and Reporting and Analytics respectively. For information about Cloud Network Automation and Reporting and Analytics, see [Deploying Cloud Network Automation](#) and [Infoblox Reporting and Analytics](#).

To delete an admin group:

1. From the **Administration** tab, select the **Administrators** tab -> **Groups** tab -> *admin\_group* checkbox, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.  
To delete an admin role:
3. From the **Administration** tab, select the **Administrators** tab -> **Roles** tab -> *admin\_role* checkbox, and then click the Delete icon.
4. In the *Delete Confirmation* dialog box, click **Yes**.

## Configuring Password Duration for Admin Groups

You can set a time duration for the password for each admin group such that the password is valid only for that duration. After the specified duration expires, the password for the users of the group expires.



#### Warning

The password expiry settings are applicable only to local users.

To set the time duration for a password for each admin group:

1. Go to the **Administration** tab, **Administrators** tab -> **Groups** tab, and select the checkbox next to the group for which you want to set the password time duration, and then click the **Edit** icon.
2. Click the **Password** tab.
3. Click the **Override** button if you want the time duration that specify here to override the time duration you set when [specifying admin passwords using Grid Properties Editor](#).  
Note that the options in the screen are enabled only if you click the **Override** button. If you do not click **Override**, the time duration you set when [specifying admin passwords using Grid Properties Editor](#) applies.
4. Select the **Password must expire** checkbox.
5. In the **Password must expire every \_ days** field, enter the number of days for which the password must be valid. For example, if you enter 11, the password is valid for 11 days.
6. In the **Reminder \_ days prior to expiration** field, enter the number of days before the expiry that NIOS sends a reminder. The range of days is from 1 to 30. The number that you enter here must always be lower than the number you enter in the **Password must expire every \_ days** field.
7. Click **Save & Close**.



#### Note

- If you click the **Override** button and do not select the **Password must expire** checkbox, it means that the password for the admin group will never expire.
- The time duration that you set here does not apply to the `saml_group` and `splunk-reporting` groups.

## Viewing Admin Groups

You can view the list of admin groups that are currently in the Grid. To view admin groups, from the **Administration** tab, select the **Administrators** tab -> **Groups** tab.

Grid Manager displays the following information:

- **Name:** The name of the admin group.
- **Superuser:** Indicates whether the admin accounts that you assign to this group have full authority to view and configure all types of data. The value can be **Yes** or **No**.
- **Comment:** The information about the admin group.  
You can select the additional fields, **Disabled** and **Site**, for display.  
You can also do the following:
  - Sort the data in ascending or descending order by column.
  - Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
  - Create a quick filter to save frequently used filter criteria. For information about using quick filters, see [Finding and Restoring Data](#).
  - Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [About the Grid Manager Interface](#).
- Print or export the data.

## Viewing Admin Roles

You can view the list of admin roles that are currently in the Grid. To view admin roles, from the **Administration** tab, select the **Administrators** tab -> **Roles** tab.

Grid Manager displays the following information:

- **Name:** The name of the admin role.
- **System:** Indicates whether the admin role is system defined or not. The value can be **Yes** or **No**.
- **Comment:** The information about the admin role.

You can select the additional fields, **Disabled** and **Site**, for display. You can also do the following:

- Sort the data in ascending or descending order by column.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information about using quick filters, see [Finding and Restoring Data](#).
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [About the Grid Manager Interface](#).
- Print or export the data.

## Viewing Admin Group Assignments

After you define permissions for an admin role, you can assign it to multiple admin groups. You can view the list of admin groups to which an admin role is assigned, as follows:

1. From the **Administration** tab, select the **Administrators** tab -> **Roles** tab -> *admin\_group* checkbox, and then click the Edit icon.
2. In the *Role* editor, select the **Admin Groups** tab.

Grid Manager displays the list of admin groups to which the role is assigned.

## Disabling Multiple Login Sessions

You can disallow multiple logins for the same NIOS session. That is, if one user in the group has logged on to a NIOS session, for example <https://255.255.255.0>, no other users in the group can log on to the same IP address from another browser or from another system.



### Warning

Disabling multiple login sessions is possible only for local users.

To do this:

1. Go to the **Administration** tab, **Administrators** tab -> **Groups** tab, and select the checkbox next to the group for which you want to disallow multiple logins and click the Edit icon.
2. Click the **Security** tab.
3. Click the **Override** button if you want to override the multiple login sessions setting that you [specified for the Grid](#).
4. Select the **Disable Concurrent Login** checkbox to disallow a member of the group to log on to multiple sessions of the same NIOS system; that is to disallow multiple login sessions per user.  
Note that Before you disable multiple logins for a group in a NIOS system, ensure that all existing sessions (if any) of members of that group in that NIOS system are logged out. If not, the existing sessions will continue to remain active even after you disable multiple logins.
5. Click **Save & Close**.

## Disabling Inactive Users

You can disable a group of users who have not logged in to NIOS for a specified duration of time.



#### Warning

Disabling inactive users is possible only for local users.

To do this:

1. Go to the **Administration** tab, **Administrators** tab -> **Groups** tab, and select the checkbox next to the group for which you want to disable users.
2. Click the **Security** tab.
3. Click the **Override** button if you want to override the disable setting that you *specified for the Grid*.
4. Select the **Disable Inactive Users** checkbox.
5. In the **Disable account if user has not logged in for <time period> days** field, specify the time period (in days) after which users who have not logged in must be disabled. The range of days is from 2 to 9999. You can also specify a reminder to be sent in the **Remind <days> prior to expiration** field. The range of days is from 1 to 30. The number of days you specify in this field is the time from which users start getting daily email reminders that their account will be disabled. NIOS sends the email reminder only if an email address has been configured for the user.
6. Select the **Allow user to reactivate account via serial console** and **Allow user to reactivate account via remote console** checkboxes if you want users to activate their account after it has been disabled. To reactivate using the serial console, see *Deploying a Single Independent Appliance*. To reactivate using the remote console, type `ssh <user name>@<ip address>`.  
Note that reactivating the account using the serial console or the remote console is possible only for superusers.
7. Click **Save & Close**.

## About Administrative Permissions

You can assign permissions to admin roles which you then assign to admin groups, or you can assign permissions directly to an admin group. The following are permissions you can grant admin groups and roles:

- **Read/Write (RW):** Allows admins to add, modify, delete, view, and search for a resource.
- **Read-Only (RO):** Allows admins to view and search for a resource. Admins cannot add, modify, or delete the resource.
- **Deny:** Prevents admins from adding, modifying, deleting, and viewing a resource. This is the default permission level for all resources.

By default, the superuser group (**admin-group**) has full access to all resources on the appliance. Superusers can create limited-access admin groups and grant them permissions to resources at the global and object levels. Limited-access admin groups must have either read-only or read/write permissions assigned in order to view information or perform tasks on any supported objects.

When you assign permissions at the global level, the permissions apply to all objects that belong to the specified resource. For example, when you define a read/write permission to all DHCP networks, the permission applies to all DHCP ranges and fixed addresses in the networks. For information about global permissions, see *Defining Global Permissions* below.

You can also define permissions at a more granular level, such as for a specific Grid member, DNS zone, Response Policy Zone, network, and even an individual database object, such as a resource record or fixed address. When you define a permission at the object level, admins with this permission can only manage the specified object and its associated objects. For information about object permissions, see *Defining Object Permissions* below.

You can use global and object permissions to restrict admins to specific DNS and DHCP resources on specific Grid members by assigning the appropriate permissions. You can use this feature to separate DNS and DHCP administration on selected Grid members. For more information, see *Defining DNS and DHCP Permissions on Grid Members* below. You can configure global permissions, object permissions, and member DNS and DHCP permissions for default and custom admin groups and roles. You cannot however define permissions for the factory default roles, such as DHCP Admin.

The appliance supports the following permissions:

Permissions	Description
Grid permissions	Includes Grid DNS properties, Grid DHCP properties, all Grid members, Microsoft servers that are managed by the Grid, network discovery, task scheduling, CSV imports, and all dashboard tasks.
IPAM permissions	Includes network views, IPv4 and IPv6 networks, and host records.
DHCP permissions	Includes Grid DHCP properties, network views, IPv4 networks, host records, DHCP ranges, DHCP fixed addresses/reservations, DHCP enabled host addresses, Mac filters, shared networks, DHCP templates, lease history, and roaming hosts.
DNS permissions	Includes Grid DNS properties, DNS views, DNS zones, Response Policy Zones, host records, bulk hosts, all DNS resource records, all shared records, and adding a blank A/AAAA record.
File distribution permissions	Includes Grid-level file distribution properties.
Reporting permissions	Includes Grid-level reporting properties.
Administration permissions	Includes all certificate authentication services, CA certificates and object change tracking.
GLB (Global Load Balancer) permissions	Includes all NIOS managed GLB objects.
DHCP fingerprint permissions	Includes all DHCP fingerprint related objects.
Named ACL permissions	Includes all named ACLs (access control lists).
Cloud permissions	Includes all tenant objects.
Super Host Permissions	Includes all super host objects.

NIOS applies permissions hierarchically in a parent-child structure. When you define permissions for a resource, all objects within that resource inherit the same permissions. For example, when you grant an admin group read/write permission for a network, the admin group automatically has read/write permission for objects in that network. To override permissions set at a higher level, you define permissions at a more specific level. For example, you can override the read/write network-level permission by setting read-only or deny permission for a fixed address or a DHCP-enabled host address. To define permissions for a more specific level, see the following:

- Permissions for common tasks, as described in [Administrative Permissions for Common Tasks](#).
- Permissions for the Grid and Grid members, as described in [Administrative Permission for the Grid](#).
- Permissions for IPAM resources, such as IPv6 networks, as described in [Administrative Permissions for IPAM Resources](#).
- Permissions for DNS resources, such as DNS views and A records, as described in [Administrative Permissions for DNS Resources](#).
- Permissions for DNS resource with associated IP addresses in networks and ranges, as described in [Administrative Permissions for DNS Resources with Associated IP addresses in Networks and Ranges](#).



- Permissions for DHCP resources, such as network views and fixed addresses, as described in [Administrative Permissions for DHCP Resources](#).
- Permissions for file distribution services, as described in [Administrative Permissions for File Distribution Services](#).
- Permissions for certificate authentication services and CA certificates, as described in [Administrative Permissions for Certificate Authentication Services and CA Certificates](#).
- Permissions for object change tracking, as described in [Administrative Permissions for Object Change Tracking](#)
- Permissions for GLB and GLB objects, as described in [Administrative Permissions for Load Balancers](#).
- Permissions for Cloud objects, as described in [Administrative Permissions for Cloud Objects](#).

When you set permissions that overlap with existing permissions, Grid Manager displays a warning about the overlaps. You can view detailed information and find out which permissions the appliance uses and which ones it ignores. For information, see [Applying Permissions and Managing Overlaps](#) below.

## Defining Global Permissions

You can define permissions at a global level for an admin group or admin role. To define global permissions:

1. For an admin group: From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin\_group* in the Groups table, and then click the Add icon -> **Global Permissions** from the Create New Permission area or select Add -> **Global Permissions** from the Toolbar.  
or  
For an admin role: From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin\_role* in the Roles table, and then click Add icon -> **Global Permissions** from the Create New Permission area or select Add -> **Global Permissions** from the Toolbar.
2. Grid Manager displays the *Manage Global Permissions* editor. For an admin group, the appliance displays the selected admin group in the **Group Permission** field. For an admin role, the appliance displays the selected admin role in the **Role Permission** field. You can also select a different group or role from the drop-down list.
3. Select the resources that you want to configure from the **Permission Type** drop-down list. Depending on your selection, Grid Manager displays the corresponding resources for the selected permission type in the table.
4. Select **Read/Write**, **Read-Only**, or **Deny** for the resources you want to configure. By default, the appliance denies access to resources if you do not specifically configure them.
5. Optionally, select additional resources from the **Permission Type** drop-down list. Grid Manager appends the new resources to the ones that you have already configured. Define the permissions for the resources you select.
6. Save the configuration and click **Restart** if it appears at the top of the screen.

The below Global Permissions table lists global permissions you can assign to admin groups or admin roles:

### Global Permissions

Permissions (Read/Write, Read-Only, or Deny)		
Administration Permissions	All Certificate Authentication Services	For more information, see <a href="#">Administrative Permissions for Certificate Authentication Services and CA Certificates</a> .
	All CA Certificates	
	Object Change Tracking	For more information, see <a href="#">Administrative Permissions for Object Change Tracking</a> .
Cloud Permissions	All Tenants	For more information, see <a href="#">Administrative Permissions for Cloud Objects</a> .



Permissions (Read/Write, Read-Only, or Deny)		
Named ACL Permissions	Named ACL	For more information, see <a href="#">Administrative Permissions for Named ACLs</a> .
DHCP Permissions	Grid DHCP Properties	For more information, see <a href="#">Administrative Permissions for Common Tasks</a> .
	All Network Views	For more information, see <a href="#">Administrative Permissions for DHCP Resources</a> .
	All IPv4/IPv6 Networks	For more information, see <a href="#">Administrative Permissions for IPv4 and IPv6 Networks and Shared Networks</a> .
	All Hosts	For more information, see <a href="#">Administrative Permissions for IPAM Resources</a> .
	All DHCP Fingerprints	For more information, see <a href="#">Administrative Permissions</a> .
	All DHCP MAC Filters	For more information, see <a href="#">Administrative Permissions for DHCP Resources</a> .
	All IPv4/IPv6 DHCP Fixed Addresses/Reservations	For more information, see <a href="#">Administrative Permissions for IPv4 or IPv6 Fixed Addresses and IPv4 Reservations</a> .
	All IPv4/IPv6 Host Addresses	For more information, see <a href="#">Administrative Permissions for DHCP Resources</a> .
	All IPv4/IPv6 Ranges	For more information, see <a href="#">Administrative Permissions for IPv4 and IPv6 DHCP Ranges</a> .
	All IPv4/IPv6 Shared Networks	For more information, see <a href="#">Administrative Permissions for IPv4 and IPv6 Networks and Shared Networks</a> .
	All IPv4/IPv6 DHCP Templates	For more information, see <a href="#">Administrative Permissions for IPv4 or IPv6 DHCP Templates</a> .
	All Microsoft Superscopes	For more information, see <a href="#">Administrative Permissions for IPv4 or IPv6 DHCP Templates</a> .
	All Roaming Hosts	For more information, see <a href="#">Administrative Permissions for Roaming Hosts</a> .
	DHCP IPv4/IPv6 Lease History	For more information, see <a href="#">Administrative Permissions for the IPv4 and IPv6 DHCP Lease Histories</a> .
DNS Permissions Grid	DNS Properties	For more information, see <a href="#">Administrative Permissions for Common Tasks</a> .

Permissions (Read/Write, Read-Only, or Deny)		
	All DNS Views	For more information, see <a href="#">Administrative Permissions for Common Tasks</a> .
	All DNS Zones	For more information, see <a href="#">Administrative Permissions for Common Tasks</a> .
	All Hosts	For more information, see <a href="#">Administrative Permissions for Hosts</a> .
	All IPv4/IPv6 Host Addresses	For more information, see <a href="#">Administrative Permissions for DNS Resources with Associated IP addresses in Networks and Ranges</a> .
	All Resource Records (A, AAAA, CAA, CNAME, DNAME, NAPTR, MX, PTR, SRV, TXT, TLSA and Bulkhost)	For more information, see <a href="#">Administrative Permissions for Common Tasks</a> .
	All Shared Record Groups	For more information, see <a href="#">Administrative Permissions for Shared Record Groups</a> .
	All Shared Records (A, AAAA, MX, SRV and TXT)	For more information, see <a href="#">Administrative Permissions for Common Tasks</a> .
	All Rulesets (BLACK List Rulesets and NXDOMAIN Rulesets)	For more information, see <a href="#">Administrative Permissions for DHCP Resources</a> .
	All DNS64 Synthesis Groups	For more information, see <a href="#">Administrative Permissions for DNS64 Synthesis Groups</a> .
	All Response Policy Zones	For more information, see <a href="#">Administrative Permissions for Zones</a> and <a href="#">License Requirements and Admin Permissions</a> .
	All Response Policy Rules	For more information, see <a href="#">Administrative Permissions for Zones</a> and <a href="#">License Requirements and Admin Permissions</a> .
	All DTC Objects (LBDN Records, LBDNs, Pools, Servers, Monitors, Certificates, GeoIP and Topologies)	For more information, see <a href="#">Administrative Permissions for Zones</a> and <a href="#">License Requirements and Admin Permissions</a> .
	Adding a blank A/AAAA record	For more information, see <a href="#">Administrative Permissions for Common Tasks</a> .
File Distribution Permissions	Grid File Distribution Permissions	For more information, see <a href="#">Administrative Permissions for File Distribution Services</a> .

Permissions (Read/Write, Read-Only, or Deny)		
Grid Permissions	All Members	For more information, see <a href="#">Administrative Permissions for Common Tasks</a> .
	Network Discovery	For more information, see <a href="#">Administrative Permissions for Discovery</a> .
	Schedule Tasks	For more information, see <a href="#">Administrative Permissions for Scheduling Tasks</a> .
	CSV Import	For more information, see <a href="#">Administrative Permissions for Named ACLs</a> .
	All Microsoft Servers	For more information, see <a href="#">Administrative Permissions for Microsoft Servers</a> .
	All Dashboard Tasks	For more information, see <a href="#">Administrative Permissions for Dashboard Tasks</a> .
	All Kerberos keys	For more information, see <a href="#">Configuring GSS-TSIG keys</a> .
	All Active Directory Domains	For more information, see <a href="#">Managing Active Directory Sites</a> .
IPAM Permissions	All Network Views	For more information, see <a href="#">Administrative Permissions for Common Tasks</a> .
	All IPv4 Networks	For more information, see <a href="#">Administrative Permissions for IPv4 and IPv6 Networks and Shared Networks</a> .
	All IPv6 Networks	For more information, see <a href="#">Administrative Permissions for IPv4 and IPv6 Networks and Shared Networks</a> .
	All Hosts	For more information, see <a href="#">Administrative Permissions for Hosts</a> .
	All IPv4 Host Addresses	For more information, see <a href="#">Administrative Permissions for DNS Resources with Associated IP addresses in Networks and Ranges</a> .
	All IPv6 Host Addresses	For more information, see <a href="#">Administrative Permissions for DNS Resources with Associated IP addresses in Networks and Ranges</a> .
	Port Control	For more information, see <a href="#">Administrative Permissions for Discovery</a> .

Permissions (Read/Write, Read-Only, or Deny)		
SAML Permissions	SAML Authentication Services	For more information, see <a href="#">Administrative Permissions for SAML</a> .
Super Host Permissions	Super Host Permissions	For more information, see <a href="#">About Administrative Permissions for Super Hosts</a> .
Security Permissions	Grid Security Permissions	For more information, see <a href="#">Administrative Permissions</a> .
Reporting Permissions	Grid Reporting Permissions	For more information, see <a href="#">Administrative Permissions for Common Tasks</a> .
	Reporting Dashboard	For more information, see <a href="#">Administrative Permissions for Reporting</a> .
	Reporting Search	For more information, see <a href="#">Administrative Permissions for Reporting</a> .
VLAN Permissions	VLAN views, VLAN ranges, and VLAN objects	For more information, see <a href="#">Administrative Permissions for VLAN Management</a> .

## Defining Object Permissions


You can add permissions to specific objects for selected admin groups or roles. When you add permissions to objects, you can select multiple objects with the same or different object types. When you select multiple objects with the same object type, you can apply permissions to the selected objects as well as the sub object types that are contained in the selected objects. As described in the below figure *Selecting Multiple Objects with the Same Object Type*, when you select five DNS forward-mapping authoritative zones, the appliance displays the object type "AuthZone" for all the zones. Since all five DNS zones are of the same object type, you can also apply permissions to all the resource records in these zones. The appliance displays the resources in the resource section of the *Create Object Permissions* editor. You can choose one or more of the resources to which you want to apply permissions.

In Cloud Network Automation, admin groups and admin users who have cloud API access have full permissions to delegated. However, you must specifically assign permissions for objects that have not been delegated in order for any admin groups or admin users to gain permission to these objects. Therefore, an admin group that has access to the cloud API would have full permissions to all delegated objects but limited permissions to non-delegated objects.

For information about how to allow cloud API access to an admin group, see [Creating Limited-Access Admin Groups](#). For information about guidelines for authority delegation, see [About Authority Delegation](#).

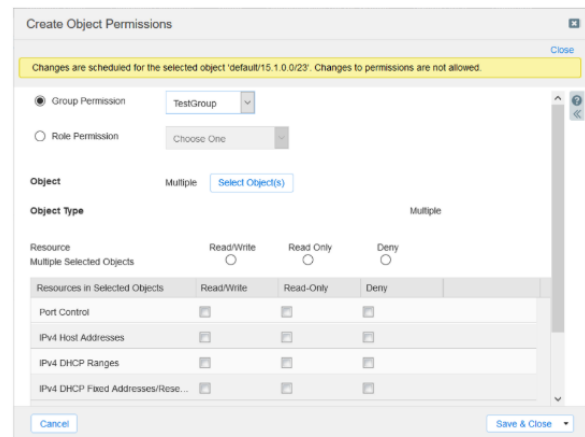
*Selecting Multiple Objects with the Same Object Type*

1 You select five forward-mapping authoritative DNS zones that have resource records such as A records, Hosts, and CNAME records.



2 Since all DNS zones have the same object type, you can apply object permissions to all the DNS zones as well as to all the resource records in the DNS zones.


3 The appliance displays the common resources in the Resources in Selected Objects column.



When you select multiple objects with more than one object type, you can add permissions to the selected objects as well as to the sub object types that are common among the selected objects. For example, when you select three DNS forward-mapping authoritative zones and two DNS IPv4 reverse-mapping authoritative zones as illustrated in the below figure Multiple Objects with Common Sub Object Types, you can apply permissions to all the five DNS zones as well as to the CNAME, DNAME, and host records in these zones because CNAME, DNAME, and host records are the common sub object types in these zones.

#### Multiple Objects with Common Sub Object Types

When you select three DNS forward-mapping authoritative zones and two IPv4 reverse-mapping authoritative zones, you can apply object permissions to all the DNS zones as well as the CNAME, DNAME and Host records in these DNS zones.



To define object permissions for an admin group or role:

- For an admin group: From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin\_group* in the Groups table, and then click the Add icon -> **Object Permissions** from the Create New Permission area or select **Add** -> **Object Permissions** from the Toolbar.  
or  
For an admin role: From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin\_role* in the Roles table, and then click Add icon -> **Object Permissions** from the Create New Permission area or select **Add** -> **Object Permissions** from the Toolbar.

2. Grid Manager displays the *Create Object Permissions* wizard. For an admin group, the appliance displays the selected group in the **Group Permission** field. For an admin role, the appliance displays the selected admin role in the **Role Permission** field. You can also select a different group or role from the drop-down list.
3. Click **Select Object(s)**. Grid Manager displays the *Object Selector* dialog box.
4. In the *Object Selector* dialog box, complete the following:
  - Enter a value or partial value of an object in the first field. This field is not case-sensitive. For example, if the object to which you want to define permissions contains "Infoblox", enter Infoblox here.
  - Select the object type for which you are searching in the **Type** drop-down list. By default, the appliance searches all object types.
  - In the operator drop-down list, select an operator for the filter criteria. Depending on what you select in the first filter field, this list displays the relevant operators for the selection.
  - In the value field, enter or select the attribute value for the first filter field. Depending on what you select for the first two filter fields, you can either enter a value or select a value from a drop-down list.
5. Click **Search**. The appliance lists all matching objects in the table. You can select multiple object types by clicking the Add icon to add more filter criteria. You can also click **Reset** to clear all entries.
6. Select the checkboxes of the objects to which you are defining permissions, and then click the Select icon.
7. In the *Create Object Permissions* wizard, do the following:
  - **Object**: Displays the name of the selected object. When you select multiple objects, the appliance displays **Multiple** here. Mouse over to the information icon to view the list of objects to which you are defining permissions.
  - **Object Type**: Displays the object type of the selected object. When you select more than one object type, the appliance displays **Multiple** here.
  - **Resource**: Displays the selected objects. When you select more than one object type, the appliance displays **Multiple Selected Objects** here. Mouse over to the information icon to view the list of objects to which you are defining permissions. Grant the resources an appropriate permission: **Read/Write**, **Read-Only**, or **Deny**.
8. Save the configuration and click **Restart** if it appears at the top of the screen.

Grid Manager displays a warning message when the permissions you define here overlap with other permissions in the system. Click **See Conflicts** to view the overlapping permissions in the *Permissions Conflict* dialog box. For information, see [Applying Permissions and Managing Overlaps](#) below.

You can also set permissions for specific objects from the objects themselves. For example, to define permissions for a particular Grid member, navigate to that Grid member and define its permissions.

To define the permissions of a specific object:

1. Navigate to the object. For example, to define permissions for a particular network, from the **Data Management** tab, select the **IPAM** tab -> *network* checkbox, and then click the Edit icon.
2. In the editor, select the **Permissions** tab, and then do one of the following:
  - Click the Add icon to add permission to the object. In the **Admin Group / Role Selector** dialog box, select an admin group or role to which you want to assign the permission, and then click the Select icon.
  - Modify the permission and resource type of a selected admin group or role.
  - Select an admin group or role and click the Delete icon to delete it.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Defining DNS and DHCP Permissions on Grid Members

You can restrict certain admin groups or roles to perform specific DNS and DHCP tasks on specific Grid members by assigning the correct global and object permissions. You can use this feature to separate the DNS and DHCP administration on different Grid members. For example, you can create an admin group or role that can only create, modify, and delete DHCP ranges in a specific network on a specific member in the Grid. This admin group or role is restricted to the specified tasks on the selected Grid member. It cannot perform other DNS or DHCP tasks on this member, and it cannot perform the specified tasks on other Grid members.

For example, you can define permissions that allow admins to create, modify, and delete DHCP ranges in network 10.0.0.0/8 on Grid member "sales.infoblox.com" by granting read/write object permissions to all DHCP ranges, network 10.0.0.0/8, and member DHCP on sales.infoblox.com. Admins with these permissions can only add, modify, and delete DHCP ranges in network 10.0.0.0/8 on Grid member sales.infoblox.com. They cannot perform other DHCP or DNS tasks on the member, and they cannot perform these tasks on other Grid members.

For information about required permissions for specific DNS and DHCP tasks, see [Administrative Permissions for](#)

### Common Tasks.

You can define the following DNS and DHCP permissions for an admin group or role:

- Grid DNS or Grid DHCP: Admins with read/write permissions can manage any DNS or DHCP resources on any Grid members. They can also modify Grid DNS or Grid DHCP properties and any member DNS and member DHCP properties. Admins with read-only permissions can only view DNS or DHCP resources. They cannot modify any DNS or DHCP resources or restart related services.
- Member DNS or Member DHCP: Admins with read/write permissions can perform the defined DNS or DHCP tasks only on the specified Grid member, not any other members. They can also modify DNS or DHCP properties on the specified member. Admins with read-only permission cannot assign the Grid member to any DNS or DHCP resources.
- Restart DNS or Restart DHCP on member: Admins with read/write permissions can restart the DNS or DHCP service on the specified Grid member, not any other members. However, they cannot modify DNS or DHCP properties on the member. They can assign the specified Grid member to any DNS or DHCP resources, but they cannot assign any other Grid members to DNS or DHCP resources.

To specify member DNS and DHCP permissions, define DNS or DHCP permissions at the global or object level for an admin group or admin role, as described in [Defining Global Permissions](#) and [Defining Object Permissions](#) above. Ensure that you include the Grid member object to which you want to restrict DNS or DHCP administration. You can assign valid permissions to administrators to manage kerberos keys. For more information, see [Configuring GSS-TSIG keys](#). You can also control whether the admins can modify DNS or DHCP properties on a member, as described in [Modifying Permissions on a Grid Member](#) below.

### Modifying Permissions on a Grid Member

Admins can perform different tasks on a Grid member based on the permissions they have. The following table Member Permissions and Tasks outlines the permissions and the tasks admins can perform on a Grid member:

#### Member Permissions and Tasks

After you add permissions to an admin group or role for a specific Grid member, you can modify the member permissions and resources. Note that when you modify the member permissions and resources, the appliance updates the permissions of the admin group or role accordingly.

	Grid Member	Member DNS or DHCP Properties	Restart DNS or DHCP on Grid Member
Read/Write	<ul style="list-style-type: none"><li>• Modify member properties</li><li>• Restart, reboot, and shutdown member</li><li>• Modify member DNS and DHCP properties</li><li>• Restart member DNS and DHCP services</li><li>• Assign and un-assign member to DNS and DHCP objects</li></ul>	<ul style="list-style-type: none"><li>• Modify member DNS or DHCP properties</li><li>• Restart member DNS or DHCP service</li><li>• Assign and un-assign member to DNS or DHCP objects</li></ul>	<ul style="list-style-type: none"><li>• Restart member DNS or DHCP service</li><li>• Assign and un-assign member to DNS or DHCP objects</li></ul>
Read-only	<ul style="list-style-type: none"><li>• View member DNS and DHCP properties</li></ul>	<ul style="list-style-type: none"><li>• View member DNS or DHCP properties</li></ul>	<ul style="list-style-type: none"><li>• N/A (You cannot define a read-only permission)</li></ul>

	Grid Member	Member DNS or DHCP Properties	Restart DNS or DHCP on Grid Member
Deny	<ul style="list-style-type: none"> <li>• Cannot modify member, DNS, and DHCP properties</li> <li>• Cannot restart related services</li> <li>• Cannot assign member to DNS and DHCP objects</li> </ul>	<ul style="list-style-type: none"> <li>• Cannot modify member, DNS, and DHCP properties</li> <li>• Cannot restart related services</li> <li>• Cannot assign member to DNS and DHCP objects</li> </ul>	<ul style="list-style-type: none"> <li>• Cannot modify member, DNS, and DHCP properties</li> <li>• Cannot restart related services</li> <li>• Cannot assign member to DNS and DHCP objects</li> </ul>

To modify Grid member permissions:

1. From the **Data Management** tab, select the **DHCP** or **DNS** tab -> **Members** tab -> *Grid\_member*, and then click the Edit icon.
2. In the *Member DHCP Properties* or *Member DNS Properties* editor, select the **Permissions** tab.
3. Click a permission in the Permissions table, select a different permission from the **Permissions** drop-down list or select a different resource from the **Resources** drop-down list. Note that when you select **Restart DNS** or **Restart DHCP**, the admins with this permission can only restart the DNS or DHCP service on the selected member. They cannot modify DNS or DHCP properties of this member.
4. Save the configuration. Note that the appliance automatically updates the permissions of the corresponding admin group or role in the **Administration** tab.

### Applying Permissions and Managing Overlaps

When an admin tries to access an object, the appliance checks the permissions of the group to which the admin belongs. Because permissions at more specific levels override those set at a higher level, the appliance checks object permissions hierarchically—from the most to the least specific. In addition, if the admin group has permissions assigned directly to it and permissions inherited from its assigned roles, the appliance checks the permissions in the following order:

1. Permissions assigned directly to the admin group.
2. Permissions inherited from admin roles in the order they are listed in the **Roles** tab of the *Admin Group* editor.

For example, an admin from the DNS1 admin group tries to access the a1.test.com A record in the test.com zone in the Infoblox default view. The appliance first checks if the DNS1 admin group has a permission defined for the a1.test.com A record. If there is none, then the appliance checks the roles assigned to DNS1. If there is no permission defined for the a1.test.com A record, the appliance continues checking for permissions in the order listed in the Permission Checking table. The appliance uses the first permission it finds.

#### Permission Checking

The appliance checks object permissions from the most to the least specific, as listed.	For each object, the appliance checks permissions in the order listed.
<ol style="list-style-type: none"> <li>1. a1.test.com A record</li> <li>2. A records in test.com</li> <li>3. test.com</li> <li>4. All zones in the default view</li> <li>5. Default view</li> <li>6. All A records</li> <li>7. All zones</li> <li>8. All DNS views</li> </ol>	<ol style="list-style-type: none"> <li>a. DNS1 admin group</li> <li>b. Role 1, Role 2, Role 3...</li> </ol>

An admin group that is assigned multiple roles and permissions can have overlaps among the different permissions. As stated earlier, the appliance uses the first permission it finds and ignores the others. For example, as shown in the below Directly-Assigned Permissions and Roles table, if an admin group has read/write permission to all A records in the test.com zone and a role assigned to it is denied permission to test.com, the appliance provides read/write access to A



records in the test.com zone, but denies access to the test.com zone and all its other resource records.  
*Directly-Assigned Permissions and Roles*

Permission assigned to the admin group	Read/Write to all A records in the test.com zone
Permission inherited from an admin role	Deny to the test.com zone
Effective permissions	Deny to the test.com zone Read/Write to all A records in test.com zone Deny to all other resource records in test.com zone

If the group has multiple roles, the appliance applies the permissions in the order the roles are listed. If there are overlaps in the permissions among the roles, the appliance uses the permission from the role that is listed first. For example, as shown in the Multiple Roles table, the first role assigned to the admin group has read-only permission to all A records in the test.com zone and the second role has read/write permission to the same records. The appliance applies the permission from the first admin role.

*Multiple Roles*

Role 1 permission	Read-only to all A records in the test.com zone
Role 2 permission	Read/Write to all A records in test.com zone Read/Write to all MX records in test.com zone
Effective permissions	Deny to the test.com zone Read-only to all A records in the test.com zone Read/Write to all MX records in test.com zone

You can check for overlapped permissions when you add permissions to roles and to admin groups, and when you assign roles to an admin group. When you create a permission that overlaps with existing permissions, Grid Manager displays a warning message and the **SeeConflicts** link on which you click to view the overlapped permissions. For information, see [Viewing Overlapping Permissions](#) below. You can also use the quick filter **Overlaps** to filter overlapped permissions, the appliance lists permissions that overlap with other permissions. If you want to change the permission the appliance uses, you must change the order in which the roles are listed or change the permissions that are directly assigned to the admin group. For information about [Creating Limited-Access Admin Groups](#), see [About Admin Groups](#).

Viewing Overlapping Permissions

When you click **See Conflicts** to view overlapping permissions, Grid Manager displays the following information in the *Permission Overlap* dialog box:

1. **Resource:** The name of the object or resource.
2. **Type:** The object type.
3. **Permission:** The permission granted. This can be Read/Write, Read-Only, or Deny.
4. **Inherited From:** Indicates the source from which the permission is inherited.
5. **Conflict Status:** Indicates whether the permission is being used or ignored. In a permission overlap, the group permission always overrides the role permission if both permissions are set at the same level (global or object). However, if the permissions are set at different levels, the permission at a more specific level overrides that set at a higher level.
6. **Role/Group Name:** The name of the admin group or admin role.

You can click the arrow key next to the resource to view the permission that is being ignored in the overlap.

## Managing Permissions

After you define permissions for an admin group and role, you can do the following:

- View the permissions, as described in [Viewing Permissions](#) below.
- Modify the permissions, as described in [Modifying Permissions](#) below.
- Delete the permission, as described in [Deleting Permissions](#) below.

## Viewing Permissions

Only superusers can view the permissions of all admin groups.

To view the permissions of an admin group or role:

1. From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab.
2. For an admin group: Select an admin group in the Groups table.  
or  
For an admin role: Select an admin role in the Roles table.
3. Grid Manager displays the following information in the Permissions table:
  - **Group/Role**: The name of the admin group or role.
  - **Permission Type**: The type of permissions. This can be Administration Permissions, Analytics Permissions, Cloud Permissions, Named ACL Permissions, DHCP Permissions, DNS Permissions, File Distribution Permissions, Grid Permissions, IPAM Permissions, Reporting Permissions, or Security Permissions.
  - **Resource**: The name of the object. For example, this field displays **All Hosts** if you have defined permissions for all the hosts in the Grid.
  - **Resource Type**: The object type. For example, this can be Host, PTR record, or Shared Network.
  - **Permission**: The defined permission for the resource.

When you click **Show All** for Admins, Groups, and Roles, Grid Manager displays all the admin accounts, admin groups, and admin roles in their respective tables.

## Filtering the List of Permissions

You can filter the permissions you want to view by selecting one of the following from the quick filter menu:

- **Effective Permissions**: Select to view only the permissions that the appliance is using for this group. The permissions that were ignored due to overlaps are not listed in this view.
- **Overlaps**: Select to view only the overlapped permissions.
- **All Configured Permissions**: Select to view all permissions.

## Modifying Permissions

You can modify the permissions of user-defined admin roles and admin groups. You cannot modify the permissions of system-defined admin roles. When you change the permissions of a role that has been assigned to multiple admin groups, the appliance automatically applies the change to the role in all admin groups to which it is assigned.

To modify the existing permissions of a role or an admin group:

1. From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab.
2. For an admin group: Select an admin group in the Groups table. or
3. For an admin role: Select an admin role in the Roles table.
4. In the Permissions table, select the resource that you want to modify, and then click the Edit icon.
5. In the *Manage Global Permissions* or *Create Object permissions* editor, select the new permission: **Read/Write**, **Read-Only** or **Deny** for the resource.
6. Save the configuration and click **Restart** if it appears at the top of the screen.

## Deleting Permissions

You can remove permissions from user-defined admin roles and admin groups. You cannot remove permissions from system-defined admin roles. When you remove permissions from a role, they are removed from the role in all admin groups to which the role is assigned. You can remove a permission from a group as long as it is not inherited from a role.

You cannot remove permissions that are inherited from a role.

To delete a permission:

1. From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab.
2. For an admin group: Select an admin group in the Groups table.  
or  
For an admin role: Select an admin role in the Roles table.
3. In the Permissions table, select the resource that you want to modify, and then click the Delete icon.
4. In the *Delete Permission Confirmation* dialog box, click **Yes**.

## Authenticating Administrators

The NIOS appliance supports the following authentication methods: local database, RADIUS, Active Directory, LDAP, TACACS+, and SAML. The appliance can use any combination of these authentication methods. It authenticates admins against its local database by default. Therefore, if you want to use local authentication only, you must configure the admin groups and add the local admin accounts, as described in [Creating Local Admins](#).

Depending on where admin user credentials are stored, you can configure the NIOS appliance to authenticate admins locally or remotely or using SAML. When you configure the authentication type as "local," NIOS authenticates admins against its local database. When you configure the authentication type as "remote," NIOS authenticates admins whose user credentials are stored remotely on authentication servers, such as RADIUS servers, AD domain controllers, LDAP servers, or TACACS+ servers. When you configure the authentication type as "SAML Only," NIOS authenticates admins against their user credentials in the IDP (Identity Provider).

Note the following when you configure remote authentication type for local admins:

- You cannot define two local admins that have the same name and belong to different authentication server groups.
- Only superusers can modify the authentication type for other admin accounts.
- At least one superuser account must use the remote authentication type.

To authenticate admins using RADIUS, Active Directory, TACACS+, or LDAP in addition to local authentication, you must define those services on the appliance and define the admin authentication policy. For information, see [About Remote Admins](#). To authenticate admins using SAML, see [Authenticating Admins Using SAML](#).

NIOS also supports two-factor authentication where it authenticates the following:

1. Administrators through the admin authentication policy.
2. Admin client certificates through the certificate authentication service.

For more information about two-factor authentication and how to configure it, see [Defining the Authentication Policy](#).



### Note

If you are using remote authentication, you must always have at least one local admin in a local admin group to ensure connectivity to the NIOS appliance in case the remote servers become unreachable.

## Authenticating Admins Using Two-Factor Authentication

You can configure NIOS to use the two-factor authentication method to authenticate users based on X.509 client certificates. In two-factor authentication, NIOS first negotiates SSL/TLS client authentication to validate client certificates. It then authenticates the admins based on the configured authentication policy. You must first configure an authentication policy, and then configure and enable the certificate authentication service for the two-factor authentication to take effect. NIOS uses certificate authentication service as the authentication policy. For information about how to set up an authentication policy, see [Defining the Authentication Policy](#).

Using the certificate authentication service, you can choose how the client certificate associates with the CA certificate. NIOS allows you to associate the client certificate manually and automatically. With manual certificate binding option, you must associate a certificate for a particular user manually, which is verified with the CA certificate. With automatic match

policy, NIOS extracts the username from the client certificate, which is then matched with the certificate authentication service. When you configure certificate authentication service, NIOS searches the CA certificates associated with each admin group to detect a valid certificate authentication service for the client's certificate. You can either select a direct match or an automatic match for a certificate authentication service.

The Infoblox certificate authentication service uses the OCSP, which is an internet protocol that validates certificate status for X.509 digital certificates that are assigned to specific admins. NIOS allows you to choose Authority Information Access (AIA) extension from a certificate as a source of OCSP configuration or define OCSP servers manually. You can also disable OCSP check for a particular certificate authentication service. For more information about OCSP, refer to RFC 2560 at <https://tools.ietf.org/html/rfc2560>.

The status of these client certificates is stored on OCSP responders to which NIOS sends requests about certificate status. A certificate status can be "good," revoked," or "unknown." After a successful SSL/TLS client authentication, NIOS authenticates the admin based on the configured authentication policy. If the authentication fails at this point, the appliance denies access to the admin. If the authentication policy has passed, the appliance sends a request to the OCSP responder for client certificate status about the admin. If the appliance receives a "good" status from the OCSP responder, the two-factor authentication is successful. The admin can now access the appliance. If the appliance receives a "revoked" or "unknown" status from the OCSP responder, the two-factor authentication fails. The admin cannot access the appliance even though the admin authentication policy has passed.

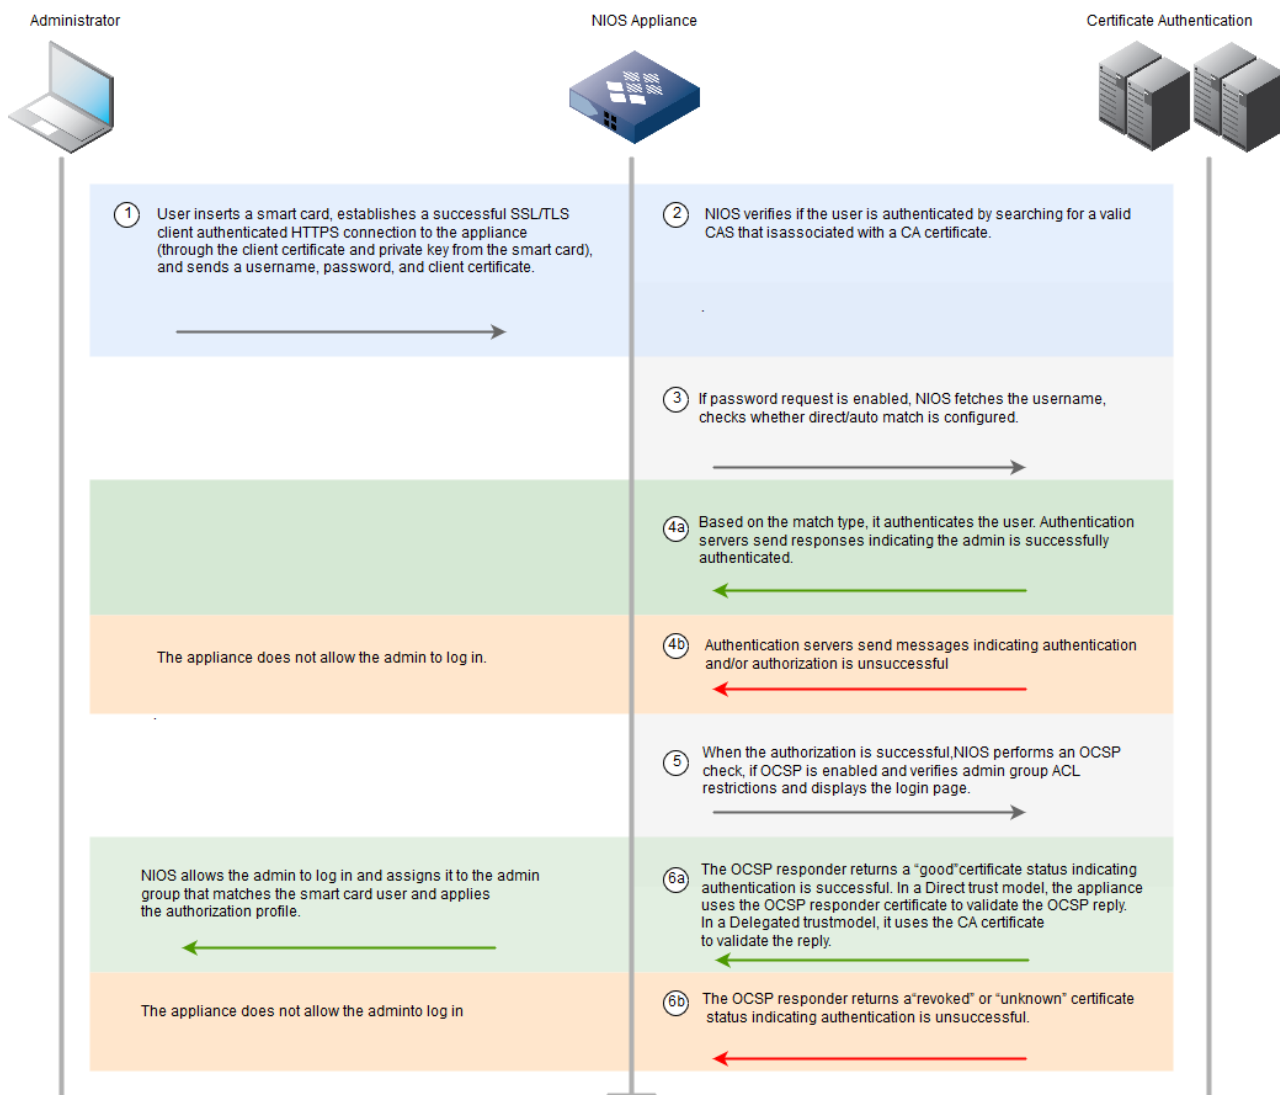
When there are multiple OCSP responders configured, the appliance contacts the responders based on their configured order. For the same client certificate, the appliance always takes the status reported by the first responder on the list that actually responds, even when there are different OCSP replies from different responders. When the appliance cannot contact the first responder or if the first responder does not reply, the appliance then takes the OCSP reply from the second responder and so on.



#### Note

- Authentication for both the admin authentication policy and OCSP validation must be successful on NIOS.
- Certificate-based authentication does not work on Cloud Platform members for WAPI calls.

*Figure 4.7 illustrates the two-factor authentication and authorization process.*



## Best Practices for Configuring Two-Factor Authentication

Only superusers and limited-access users with the correct permissions can configure two-factor authentication. For information about admin roles and permissions, see [Managing Admin Groups and Admin Roles](#). To configure two-factor authentication, consider the following:

- You must first set up a certificate authentication service and enable it.
- You can configure only one certificate authentication service that contains one or multiple OSCP responders to which NIOS sends requests about client certificate status. The appliance supports IPv4 and IPv6 OSCP responders.
- When you configure multiple OSCP responders, you can put them in an ordered list. The appliance contacts the first responder on the list. If the connection fails, it moves on to the second one, and so on. The result of the status check for a client certificate is based on the status reported by the first responder that replies.
- You can configure the timeout value and retry attempts that the appliance waits and tries before it moves on to the next OSCP responder.
- You can upload server certificates for each responder for OSCP response validation. You must upload an OSCP server certificate if you select the direct trust model.
- You can disable a specific responder if the server is out of service for a short period of time.
- Before you add an OSCP responder to the server group, you can test the server credentials.

To configure and enable two-factor authentication, complete the following tasks:

1. For local and remote authentication, ensure that the admin names for smart card users match the CNs (Common Names) used in the client certificates. For information about local and remote authentication, see [About Admin Accounts](#).
2. Upload the CA (Certificate Authority) certificate, as described in About CA Certificates, see [Managing Certificates](#). The CA-signed certificates are used to validate OCSP server certificates and admin OCSP client certificates. Ensure that the CA certificate is in .PEM format. The .PEM file can contain more than one certificate.

Note that the uploaded CA certificates must be the ones that issued the client certificates to be authenticated. Otherwise, clients such as browsers, cannot establish a successful SSL/TLS client authenticated HTTPS session to the appliance.

3. Configure a certificate authentication service and enable it, as described in [Configuring Certificate Authentication Services](#) below.
4. View certificate authentication services, as described in [Viewing Certificate Authentication Services](#) below.
5. Modify certificate authentication services, as described in [Modifying Certificate Authentication Services](#) below.
6. Delete certificate authentication services, as described in [Deleting Certificate Authentication Services](#) below.

Note that once you save the certificate authentication service configuration, the appliance terminates administrative sessions for all admin users. After you enable the certificate authentication service, you can verify whether two-factor authentication is enabled. Go to the **Administration** -> **Administrators** -> **Authentication Policy** tab, Grid Manager displays the "Two-Factor Authentication Enabled" banner in this tab.

## Configuring Certificate Authentication Services

To configure and enable the certificate authentication service, complete the following:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the **Certificate Authentication Services** subtab and click the Add icon.
3. In the *Add Certificate Authentication Service* wizard, complete the following:
  - **Name:** Enter a name for the certificate authentication service.
  - **Username/password request:** Select the checkbox if the certificate authentication service must request username and password from the user. When you select this checkbox, NIOS populates the username from the certificate and requests password from the user. If you do not select the checkbox, only the certificate is necessary to log in to the appliance. The appliance ignores the username and password when the user provides both. You can only see the login button and do not have to provide the password. The appliance displays the username when you click the login button.
  - **Auto-populate username:** Select a value from the drop-down list. You can define how the appliance must authenticate a particular user and its associated group. The values in the list are **Auto-match** and **Direct-match**. When you select **Direct-match**, NIOS searches for users with directly assigned certificates, which contains issuer details and serial attributes, in the local database. Users with directly assigned certificates can use certificate based authentication only.
  - **Auto match by:** Select a value from the drop-down list. This field is enabled only when you select **Auto-match** for **Auto-populate username**. NIOS extracts the username from the certificate and searches for it in effective authorization policies based on the configured match policies. The values in the list are:
    - **AD Issuer Subject:** Select this from the drop-down list to authenticate the user based on the Active Directory mentioned by the user.
    - **SAN Email:** Select this from the drop-down list to authenticate the user based on the email address in the SAN (Subject Alternative Name).
    - **SAN UPN:** Select this from the drop-down list to authenticate the user based on the UPN (User Principal Name) in the SAN (Subject Alternative Name).
    - **Serial Number:** Select this from the drop-down list to authenticate the user based on the serial number.
    - **Subject DN Common Name:** Select this from the drop-down list to authenticate the user based on the subject DN (Distinguished Name) common name. A Subject DN can include information about the user who is being authenticated, including common name, name of the organization, country code, and so on.
    - **Subject DN Email:** Select this from the drop-down list to authenticate a user based on the subject DN email address.

- **Enable remote lookup for user membership:** Select the checkbox to enable lookup on remote servers. NIOS performs lookup against local users by default. For a remote lookup, you must specify the username and password for the authentication service. You can perform a look up for a user membership only if the remote service admin that is configured for remote lookup has enough permissions to read other user's membership information. You must also select the remote service that must be used for lookup. Note that NIOS supports remote lookup for Active Directories only. Note that You can select the above checkbox, **Authentication Service** and **Service Account Credentials** fields only when you select **Auto-match** for **Auto-populate username**. You must not select the **Username/password request** checkbox when you select the checkbox for **Enable remote lookup for user membership**.
- **Authentication Service:** Select an authentication service from the drop-down list.
- **Service Account Credentials:** Enter a username and password for authenticating lookup on remote servers.
- **Comment:** Optionally, enter additional information about the certificate authentication service.
- **Disable:** Select this checkbox to disable the record. Clear the checkbox to enable it.

4. Click **Next** to save the configuration and add OCSP responders to the table.

5. You can add multiple OCSP responders for failover purposes.

- **OCSP Check Type:** Select a value from the drop-down list to perform OCSP checks. The values in the drop-down list are:
  - **AIA and Manual:** Select this from the drop-down list to use AIA (Authority Information Access) extension of X.509 certificate, when it is present, to authenticate the user. Note that AIA points to the certificate authentication service that is used to verify the certificate. If AIA is not available, then the authentication fails. If the certificate does not contain AIA, then the appliance uses manual OCSP for authentication.
  - **AIA only:** Select this from the drop-down list to use AIA only to authenticate the user. AIA points to the certificate authentication service that is used to verify the certificate. By selecting this option you restrict NIOS to use AIA only. If the certificate does not contain AIA or it is not complete, then the authentication fails.
  - **Disabled:** Select this from the drop-down list if you do not want to perform an OCSP check.
  - **Manual:** Select this from the drop-down list to define OCSP settings and upload CA certificates manually. When you select this option, NIOS ignores AIA even though it is present.
- **OCSP Responders:** Click the Add icon and complete the following in the **Add OCSP Responder** section:
  - **Server Name or IP Address:** Enter the FQDN or the IP address of the OCSP responder that is used for authentication. The appliance supports IPv4 and IPv6 OCSP responders.
  - **Comment:** Enter useful information about the OCSP responder.
  - **Port:** Enter the port number on the OCSP responder to which the appliance sends authentication requests. The default is 80.
  - **Server Certificate:** Click **Select** to upload a server certificate. In the *Upload* dialog box, click **Select** to navigate to the certificate, and then click **Upload**. The appliance validates the certificate when you save the configuration. A server certificate is required for the direct trust model.
  - **Disable Server:** Select this checkbox to disable the OCSP responder if, for example, the connection to the server is down and you want to stop the NIOS appliance from trying to connect to this server. Note that you cannot save the OCSP configuration when you disable all OCSP responders, thus the certificate authentication service is disabled and two-factor authentication is no longer in effect. You cannot add OCSP responders when you select **AIA only** or **Disabled** from the drop-down list for **OCSP Check Type**.

Click **Add** to save the configuration and add the responder to the table. You can add multiple OCSP responders for failover purposes. You can use the up and down arrows to place the responders in the order you desire. The appliance tries to connect with the first responder on the list. If the connection fails, it tries the next responder on the list, and so on. Grid Manager displays the following for each responder:

- **Responder:** The FQDN or the IP address of the OCSP responder.
- **Comment:** Information you entered about the OCSP responder.
- **Port:** The port number on the OCSP responder to which the appliance sends authentication requests.
- **Disabled:** Indicates whether the OCSP responder is disabled or not. Note that you must enable at least one responder to enable the certificate authentication service.



You can also click **Test** to test the configuration. If the appliance connects to the responder using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the responder, the appliance displays a message indicating an error in the configuration.

- **Response Timeout (s):** Enter the time the appliance waits for a response from the specified OCSP responder.

The default is 1 second. You can select the time unit from the drop-down list.

- **Retries:** Enter the number of times the appliance tries to connect to the responders after a failed attempt. The default is 5.
- **Recovery Interval:** Enter the time the appliance waits to recover from the last failed attempt in connecting to an OCSP responder. Select the time unit from the drop-down list. The default is 30 seconds. This is the time interval that NIOS waits before it tries to contact the responder again since the last attempt when the appliance could not connect with the responder or when the responder did not send a reply within the configured response timeouts and retry attempts.
- **Trust Model:** Select **Direct** or **Delegated** from the drop-down list as the trust model for OCSP responses. In a direct trust model, OCSP responses are signed with an explicitly trusted OCSP responder certificate. You must upload the OCSP responder certificate if you select **Direct**. In a delegated trust model, OCSP responses are signed with a trusted CA certificate. A server certificate is not required when you select **Delegated**. The default is **Direct**.

6. Click **Next** to save the configuration and associate CA Certificates with the respective certificate authentication service. You can associate multiple CA certificates with the service.

Note that enabling the certificate authentication service terminates administrative services for all users. Ensure that you have uploaded the correct CA certificates before enabling the service. Your login names must also match the common name used in the certificate. When you configure multiple OCSP responders, ensure that you place them in the correct order because the status check for a client certificate is based on the OCSP reply sent by the first OCSP responder that replies.

NIOS detects a valid certificate authentication service for a client's certificate by searching through the assigned CA certificates for each group. NIOS matches issuer field in the client's certificate with the CA certificate to find the appropriate match. Note that the subject in CA certificate must match the issuer in the client's certificate and corresponding certificate authentication service.

Note the following about the certificate authentication service:

- You cannot assign the same CA certificate to the same group twice or to a different certificate authentication service. However, different certificate authentication services can contain CA certificates with the same subject. To distinguish such groups you can use **Client Subject name** to determine which certificate must match the CA certificate to be associated with the certificate authentication service. If the client certificate does not match any certificate authentication service, then the authentication fails. A CA certificate verifies the client certificate.

7. Click **Add** to associate CA certificates with the certificate authentication service. The following information is displayed when you associate a CA certificate:

- **Subject:** The name of the certificate.
- **Issuer:** The name of the trusted CA that issued the certificate.
- **Valid From:** The date from which the certificate becomes valid.
- **Valid To:** The date until which the certificate is valid. You can do the following:
  - Select a certificate and click the Delete icon to delete it.
  - Print the data or export it in .csv format.

You can also do the following for a certificate authentication service:

- Use **Global Search** to search for certificate authentication services. For information about Global search, see [About the Grid Manager Interface](#).
- View audit log entries for the certificate authentication service. For information about viewing the audit log, see [Monitoring Tools](#).
- Select a certificate authentication service and click the Delete icon to delete it. In the *Delete Confirmation* dialog box, click **Yes** to confirm deletion.
- Modify a certificate authentication service as mentioned in Modifying Certificate Authentication Services below.
- Print the data or export it in .csv format.



## Enabling Certificate Authentication Service for a User

You can restrict users to use certificate based authentication only. Note that certificate authentication service with a direct-match searches only for users with certificate authentication service enabled. Such users are successfully authenticated by the certificate authentication service using auto-match.

1. From the **Administration** tab, click the **Administrators** tab -> **Admins** tab -> *admin\_account* checkbox, and then click the Edit icon.
2. In the *Administrator* editor, click the **General** tab, and then click the **Advanced** tab.
3. In the **General Advanced** tab, complete the following:
  - **Enable Certificate Authentication:** Select this checkbox to enable certificate authentication for the selected user. You must configure certificate authentication service and associate a valid client CA certificate with the selected user. This is disabled by default.
  - **Client Certificate Number:** You can specify a client certificate number only when you select the **Enable Certificate Authentication** checkbox. This is disabled by default. Enter the serial number as mentioned in the certificate. Examples: 397F9435000100000032 (hexadecimal format), 123 (decimal format), and so on.
  - **Client CA Certificate:** You must associate a CA certificate that signs the client certificate. Click **Select** to associate a CA certificate. When you select a CA certificate from the list, NIOS displays the subject of the selected CA certificate. The *CA Certificate Selector* dialog box displays the following information about CA certificates:
    - **Issuer:** The name of the trusted CA that issued the certificate.
    - **Valid From:** The date from which the certificate becomes valid.
    - **Valid To:** The date until which the certificate is valid.
    - **Subject:** The name of the certificate.Click **OK** to select and associate the client CA certificate with the selected admin user.
4. Save the configuration.

## Viewing Certificate Authentication Services

To view the certificate authentication service, complete the following:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the **Certificate Authentication Services** subtab.
3. Grid Manager displays the following about the certificate authentication service:
  - **Name:** The name of the certificate authentication service.
  - **Comment:** Comments about the certificate authentication service.

You can also display the following column:

- **Disabled:** Indicates if the certificate authentication service is enabled or disabled.

You can do the following in this tab:

- Sort the data in ascending or descending order by column.
- Select the certificate authentication service and click the Edit icon to modify data, or click the Delete icon to delete it.
- Print and export the data in this tab.

## Modifying Certificate Authentication Services

To modify a certificate authentication service:

1. From the **Administration** tab, click the **Authentication Server Groups** tab -> **Certificate Authentication Services** subtab -> select a certificate authentication service, and then click the Edit icon.
2. The *Certificate Authentication Service* editor provides the following tabs from which you can modify data:
  - **General:** In this tab, modify certificate authentication service data, as described in *Configuring Certificate Authentication Services* above.
  - **OCSP:** Modify associated OCSP responders.
  - **CA:** Add and delete CA certificates that are associated with the certificate authentication service.
3. Save the configuration.

## Deleting Certificate Authentication Services

You can delete a certificate authentication service any time after you have created it. To delete a certificate authentication service:

1. From the **Administration** tab, click the **Authentication Server Groups** tab -> **Certificate Authentication Services** subtab, select a certificate authentication service and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.

## About Remote Admins

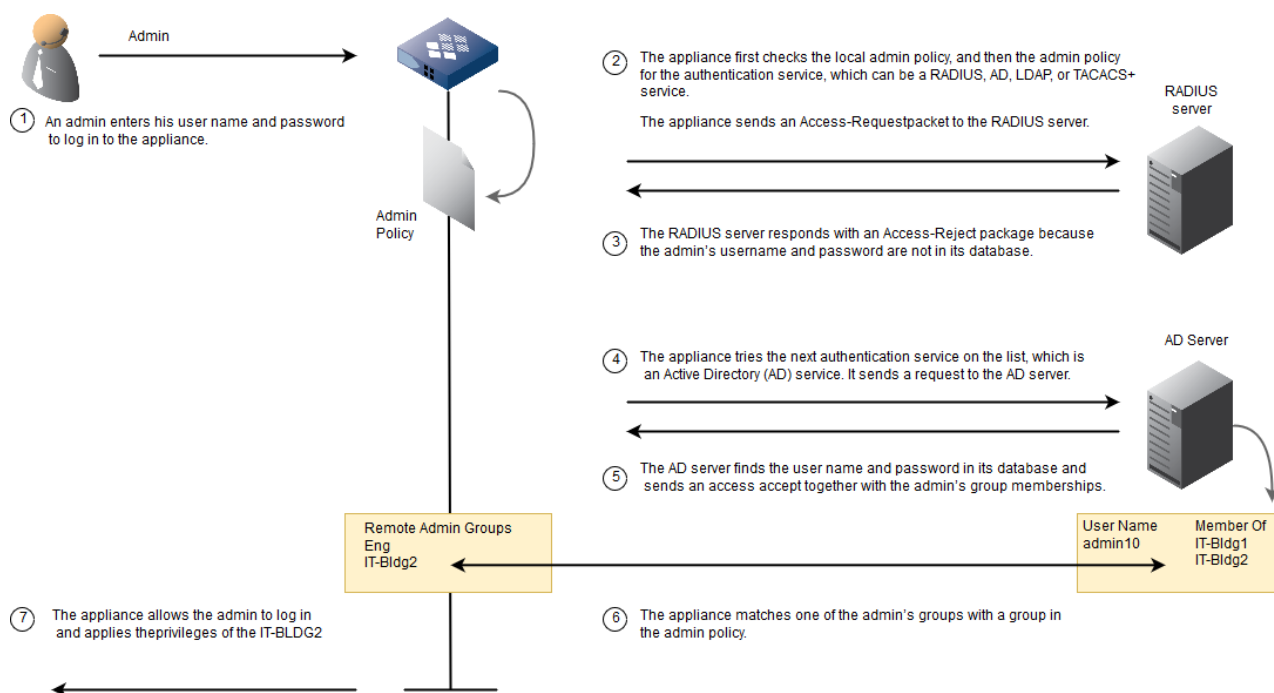
NIOS can authenticate admins whose user credentials are stored remotely on RADIUS servers, AD domain controllers, LDAP servers, or TACACS+ servers. You can configure authentication server groups for each type of server to which NIOS sends authentication requests. For example, you can create a server group for RADIUS servers and another server group for AD domain controllers. Then in the admin authentication policy, you can list which authentication server groups to use and in what order.

In addition, if admin groups are configured on the remote authentication server, you can configure admin groups with the same names on the NIOS appliance and list them in the authentication policy as well. Then if the remote authentication server provides the admin group name while authenticating an admin, NIOS can automatically assign the admin to the matching admin group specified in the authentication policy. You can also create a default admin group for all admins that are authenticated through a remote authentication service.

Managing Administrators illustrates the authentication and authorization process for remote admins. In the example, two authentication server groups are configured—a RADIUS server group and an AD server group. When an admin logs in with a user name and password, the appliance uses the first service listed in the admin policy to authenticate the admin. If authentication fails, the appliance tries the next service listed, and so on. It tries each service on the list until it is successful or all services fail. If all services fail, then the appliance denies access and generates an error message in syslog.

If authentication succeeds, the NIOS appliance tries to match the admin group names in the admin policy to any groups received from the remote server. If it finds a match, the NIOS appliance applies the privileges of that group to the admin and allows access. If the appliance does not find a match, then it applies the privileges of the default group. If no default group is defined, then the appliance denies access.

### *Authenticating Remote Admins*



Only superusers can perform the following tasks to configure NIOS to authenticate admins using remote authentication servers:

- Configure the authentication server groups. You can create multiple RADIUS, LDAP, and AD server groups, and certificate authentication services, but only one TACACS+ server group.
  - For information about RADIUS authentication, see [Authenticating Admins Using RADIUS](#).
  - For information about AD authentication, see [Authenticating Admins Using Active Directory](#).
  - For information about TACACS+ authentication, see [Authenticating Admin Accounts Using TACACS+](#).
  - For information about LDAP authentication, see [Authenticating Admins Using LDAP](#).
  - For information about the certificate authentication service, see [Authenticating Admins Using Two-Factor Authentication](#).
- Configure admin groups with names that match those on the remote server. For information about admin groups, see [About Admin Groups](#).
- Configure the admin policy, as described in [Defining the Authentication Policy](#).



#### Note

Infoblox strongly recommends that even if you are using remote authentication, you always have at least one local admin in a local admin group to ensure connectivity to the appliance in case the remote servers become unreachable. Also, when you delete an authentication server group, the appliance removes it from the system. Deleted authentication server groups are not moved to the Recycle Bin. Once deleted, the authentication server groups no longer exist in the system.

When remote authentication is successful, the appliance creates a remote admin user object in the NIOS database, which stores user preferences such as time zone, table size, and active Dashboard widgets for the remote user. If the remote user does not log in to the appliance for more than 180 days, the appliance removes the corresponding admin user object from the database. Although the remote user can still log in to the appliance, user preferences are lost. The Grid Master performs this clean up action once a day.

You can also authenticate users based on X.509 client certificates. You can configure NIOS to authenticate these admins

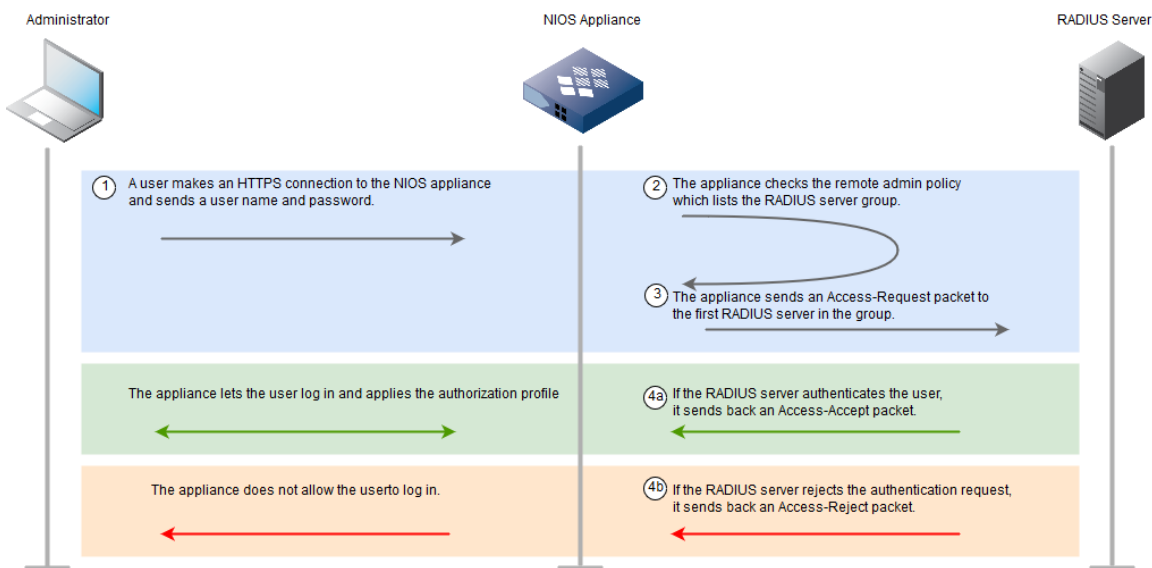
through the two-factor authentication method. For more information about two-factor authentication and how to configure it, see [Defining the Authentication Policy](#).

## Authenticating Admins Using RADIUS

RADIUS provides authentication, accounting, and authorization functions. The NIOS appliance supports authentication using the following RADIUS servers: FreeRADIUS, Microsoft, Cisco, and Funk.

When NIOS authenticates administrators against RADIUS servers, NIOS acts similarly to a network access server (NAS), which is a RADIUS client that sends authentication and accounting requests to a RADIUS server. The following figure illustrates the RADIUS authentication process.

### Authentication using a RADIUS server



## Authentication Protocols

When you configure the NIOS appliance to authenticate admins against a RADIUS server group, you must specify the authentication protocol of each RADIUS server, which can be either PAP (Password Authentication Protocol) or CHAP (Challenge Handshake Authentication Protocol).

PAP tries to establish the identity of a host using a two-way handshake. The client sends the user name and password in clear text to the NIOS appliance. The appliance uses a shared secret to encrypt the password and sends it to the RADIUS server in an Access-Request packet. The RADIUS server uses the shared secret to decrypt the password. If the decrypted password matches a password in its database, the user is successfully authenticated and allowed to log in. With CHAP, when the client tries to log in, it sends its user name and password to the NIOS appliance. The appliance then creates an MD5 hash of the password together with a random number that the appliance generates. It then sends the random number, user name, and hash to the RADIUS server in an Access-Request package. The RADIUS server takes the password that matches the user name from its database and creates its own MD5 hash of the password and random number that it received. If the hash that the RADIUS server generates matches the hash that it received from the appliance, then the user is successfully authenticated and allowed to log in.

You can configure one of the following modes to send the authentication request to the RADIUS server:

- **Ordered:** In this mode, the authentication request is sent to the first server in the list. The authentication request is sent to the next server only when the first server is out of service or unavailable.
- **Round Robin:** In this mode, the first authentication request is sent to a server chosen randomly in a group. If there is no response from the server, continued attempts are performed sequentially until it selects the last server in the list. Then it starts with the first server in the list and continues the selection process until all the servers have been attempted.

## Accounting Activities Using RADIUS

You can enable the accounting feature on the RADIUS server to track whether an administrator has initiated a session. After an administrator successfully logs in, the appliance sends an Accounting-Start packet to the RADIUS server.

## Configuring Remote RADIUS Servers

For NIOS to communicate with a RADIUS server, you must also set up the remote RADIUS server to communicate with the NIOS appliance.



### Note

If you have two Infoblox appliances in an HA pair, enter both the members of the HA pair as separate access appliances and use the LAN or MGMT IP address of both appliances (not the VIP address), if configured.

Depending on your particular RADIUS server, you can configure the following RADIUS server options to enable communication with the NIOS appliance:

- Authentication Port
- Accounting Port
- Domain Name/IP Address of the NIOS appliance
- Shared Secret Password
- Vendor Types

## Configuring Admin Groups on the Remote RADIUS Server

Infoblox supports admin accounts on one or more RADIUS servers.

On the remote RADIUS server, do the following to set up admins and associate them with an admin group:

- Import Infoblox VSAs (vendor-specific attributes) to the dictionary file on the RADIUS server
- For third-party RADIUS servers, import the Infoblox vendor file (the Infoblox vendor ID is 7779)
- Define the admin group
- Associate one or more remote admin accounts with the admin group
- Add and activate a policy for the admin accounts, but do not associate the policy with a policy group that contains an infoblox-group-info attribute.

Refer to the documentation for your RADIUS server for more information.

## Configuring RADIUS Authentication

To configure NIOS to use one or more RADIUS server groups to authenticate administrators, you must do the following:

- Configure at least one RADIUS authentication server group. For more information, see [Configuring a RADIUS Authentication Server Group](#) below.
- Define admin groups for the admins that are authenticated by the RADIUS servers and specify their privileges and settings. The group names in NIOS must match the admin group names on the RADIUS server. See [About Admin Groups](#) for information about defining admin groups.
- In the authentication policy, add the RADIUS server groups and the admin groups that match those on the RADIUS server. You can also designate an admin group as the default group for remote admins. NIOS assigns admins to this group when it does not find a matching group for a remote admin. See [Defining the Authentication Policy](#) for more information about configuring the policy.

## Configuring a RADIUS Authentication Server Group

You can add multiple RADIUS servers to the group for redundancy. When you do, the appliance tries to connect to the first RADIUS server on the list and if the server does not respond within the maximum retransmission limit, then it tries the next RADIUS server on the list. NIOS tries to connect to each RADIUS server in the order the servers are listed. If it

does not receive a response within the configured timeout period and has tried to connect the specified retry value, then it tries the next RADIUS server on the list. It logs an error to syslog when it fails to connect to any of the servers in the group.

After you add a RADIUS server to the NIOS appliance, you can validate the configuration. The appliance uses a pre-defined username and password when it tests the connection to the RADIUS server. The pre-defined user name is "Infoblox\_test\_user" and the password is "Infoblox\_test\_password". Do not use these as your administrator username and password.

To configure a RADIUS authentication server group :

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the Add icon in the **RADIUS Services** subtab.
3. In the *Add RADIUS Authentication Service* wizard, complete the following:
  - **Name:** Enter the name of the server group.
  - **RADIUS Servers:** Click the Add icon and enter the following:
    - **Server Name or IP Address:** Enter the FQDN or the IP address of the RADIUS server that is used for authentication.
    - **Comment:** Enter additional information about the RADIUS server.
    - **Authentication Port:** The destination port on the RADIUS server. The default is 1812. This field is required only if you do not enable accounting on the RADIUS server. This field is not required if you enable accounting to configure an accounting-only RADIUS server.
    - **Authentication Type:** Select either PAP or CHAP from the drop-down list. The default is PAP.
    - **Shared Secret:** Enter the shared secret that the NIOS appliance and the RADIUS server use to encrypt and decrypt their messages. This shared secret is a value that is known only to the NIOS appliance and the RADIUS server.
    - **Enable Accounting:** Select this to enable RADIUS accounting for the server so you can track an administrator's activities during a session. When you enable accounting, you must enter a valid port number in the **Accounting Port** field.
    - **Accounting Port:** The destination port on the RADIUS server. The default is 1813.
    - **Connect through Management Interface:** Select this so that the NIOS appliance uses the MGMT port for administrator authentication communications with just this RADIUS server.
    - **Disable server:** Select this to disable the RADIUS server if, for example, the connection to the server is down and you want to stop the NIOS appliance from trying to connect to this server.
    - Click **Test** to test the configuration. If the NIOS appliance connects to the RADIUS server using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the RADIUS server, the appliance displays a message indicating an error in the configuration.
    - Click **Add** to add the server to the list.

When you add multiple RADIUS servers, the appliance lists the servers in the order you added them. This list also determines the order in which the NIOS appliance attempts to contact a RADIUS server. You can move a server up or down the list by selecting it and clicking the up or down arrow.

You can also delete a RADIUS server by selecting it and clicking the Delete icon.

- **Authentication:** Optionally, modify the authentication settings. These settings apply to all RADIUS servers that you configure on the NIOS appliance.
- **Timeout(s):** Specify the number of seconds that the appliance waits for a response from the RADIUS server.
- **Retries:** Specify how many times the appliance attempts to contact an authentication RADIUS server.

The default is 5.

If you have configured multiple RADIUS servers for authentication and the NIOS appliance fails to contact the first server in the list, it tries to contact the next server, and so on.

- **Accounting:** Optionally, modify the Accounting settings.
  - **Timeout(s):** Specify the number of seconds that the appliance waits for a response from the RADIUS server.
  - **Retries:** Specify how many times the appliance attempts to contact an accounting RADIUS server. The default is 1000.
- **Mode:** Specifies how the appliance contacts the RADIUS servers. The default is Ordered List.
  - **Ordered List:** The Grid member always selects the first RADIUS server in the list when it sends an authentication request. It queries the next server only when the first server is considered down.

- **Round Robin:** The Grid member sends the first authentication request to a server chosen randomly in a group. If there is no response from the server, the Grid member selects the next server in the group. Continued attempts are performed sequentially until it selects the last server in the group. Then it starts with the first server in the group and continues the selection process until all the servers have been attempted.
- **Comment:** Enter useful information about the RADIUS service.
- **Disable:** Select this to disable RADIUS authentication for the servers listed in the table.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

Note that the following fields in the wizard do not apply to this feature: **Enable NAC Filter**, **Cache Time to Live**, and **Recovery Interval**. They are used with the NAC Integration feature described in [Authenticated DHCP](#).

## Creating Local Admins

When you create an admin account, you must specify the authentication type, name, password, and admin group of the administrator. You can also control in which time zone the appliance displays the time in the audit log and the DHCP and IPAM tabs of Grid Manager, such as the *DHCP Lease History* and *DHCP Leases* panels. The appliance can use the time zone that it automatically detects from the management system that the admin uses to log in.

Alternatively, you can override the time zone auto-detection feature and specify the time zone. To create an admin account and add it to an admin group:

- Log in as a superuser.
- From the **Administration** tab, select the **Administrators** tab -> **Admins** tab, and then click the Add icon.  
or  
From the **Administration** tab, select the **Administrators** tab -> **Groups** tab -> *admin\_group*, and then click the Add icon.
- In the *Add Administrator* wizard, complete the following:
  - **Authentication Type:** The default is **Local**. When you select **Local**, NIOS authenticates admins against the local database.
 

**Local:** The following fields are displayed when you select **Local** authentication type. Enter the following:

    - **Login:** Enter a name for the administrator. This is the user name that the administrator uses to log in to the appliance. This user name is stored in the NIOS local database.
    - **Password:** Enter a password for the administrator. This is the password that the administrator uses to log in to the appliance. This password is stored in the NIOS local database.
    - **Confirm Password:** Enter the same password.  
Note In NIOS 8.5.2 or later, when you set up the account for a Grid Master or a standalone vNIOS instance that is deployed on AWS, the minimum password length must be four characters. The password must consist of at least one uppercase character, one lowercase character, one numeric character, and one symbol character. Example: Infoblox1!  
If the symbol character is at the beginning of the password, then include the password within quotes ("). Example: '@Infoblox123'.
  - **Use AWS SSH authentication keys:** To prevent CLI login failures after upgrading, you will need to enable **Use AWS SSH authentication keys** for each user that needs CLI access to AWS appliances. When you select **Use AWS SSH authentication keys**, NIOS allows you to access the CLI either by using a key pair and entering a password, or only by using the key pair which means the password-only authentication is blocked for the user. You can upload the SSH key by using the **Manage SSH Public Keys** field. It is mandatory to upload a valid SSH public key if you select the **Use AWS SSH authentication keys** option.  
If you use the **User data** field in the AWS console to install a NIOS license, the **Use AWS SSH authentication key** option is enabled by default. For more information about the **User data** field, see the *Initializing New Infoblox vNIOS for AWS Instances with the AWS User Data Field* section in the [Installation Guide for vNIOS for AWS documentation](#).  
Note for a TE-V4025 appliance, if you use the **User data** field to install the TE-4025 license, the **Use AWS SSH authentication key** option will not be enabled by default. Therefore, Infoblox recommends that you first deploy the vNIOS instance without specifying the IB-4025 license, and then install the license from the NIOS CLI.



- **Authentication Method:** You can choose **Key pair** or **Key pair + Password** methods from the **Authentication Method** drop-down list. A server generates two distinct, but related keys: a public key that you upload and a corresponding private key that is stored in the system. A Key pair is the combination of these two related keys and is the default authentication method. If you select **Key pair** as the authentication method, then a user can access the CLI with a valid key pair. If you select **Key pair + Password** as the authentication method, the user must provide a password to access the CLI even after a successful key pair authentication. For information on defining and managing passwords, see *Managing Passwords* below.
- **Manage SSH Public Keys:** You need to upload a valid SSH public key file. The supported key types are RSA, EDSA, and ED25519. The **Key Type** and **Key Value** fields in the **MANAGE SSH PUBLIC KEYS** are automatically updated once you upload a valid SSH key.

Note that from NIOS 8.5.2 onwards, the **Use AWS SSH authentication keys**, **Authentication Method**, and **Manage SSH Public Keys** fields are not available for the **Remote** and **SAML Only** authentication types. That is, you cannot use the CLI to access vNIOS for AWS if you are a remote user or a SAML user.

- **Remote:** When you select **Remote**, NIOS authenticates admins based on the user credentials stored remotely on authentication servers, such as RADIUS servers, AD domain controllers, LDAP servers, or TACACS+ servers. The **Login** field is displayed when you select **Remote** authentication type. Enter a name for the administrator that is stored in the database of the remote server. This is the user name that the administrator uses to log in to the appliance.
- **SAML Only:** When you select **SAML Only**, NIOS authenticates admins based on the user credentials stored in the IDP (Identity Provider). An admin can log in to NIOS only by clicking the **SSO Login** button and if the user credentials exist in the IDP account.
- **SAML/Local:** When you select **SAML/Local**, NIOS authenticates admins based on the user credentials stored in the IDP, when the SSO Login button is clicked or against the local database when the User name and Password is supplied and the Login button is clicked. For **SSO Login**, the user name and password need not be supplied in the NIOS GUI, rather it should be supplied in the IDP's login prompt. For information about SAML authentication, see [Authenticating Admins Using SAML](#).



#### Note

You cannot configure the **Remote** authentication type for NIOS admin users who belong to the **fireeye-group** admin groups.

**Email Address:** Enter the email address for this administrator. The appliance uses this email address to send scheduling notifications.

- **Admin Group:** Click **Select** to specify an admin group. If there are multiple admin groups, Grid Manager displays the *Admin Group Selector* dialog box from which you can select one. An admin can belong to only one admin group at a time.

NIOS appliance creates a new group, **fireeye-group**, when you add the first FireEye zone. The FireEye admin group is read-only and you cannot assign permissions to it. Select **fireeye-group** for the admin group and add users to this group. For more information, see [About FireEye Integrated RPZs](#).



You cannot add a NIOS admin user that uses the Remote authentication type to the **fireeye-group** admin group.

- **Comment:** Enter useful information about the administrator.
- **Disable:** Select this checkbox to retain an inactive profile for this administrator in the configuration. For example, you might want to define a profile for a recently hired administrator who has not yet started work. Then when he or she does start, you simply need to clear this checkbox to activate the profile.
- **Status:** Displays the status of the administrator. The status can be one of the following:
  - **Active:** The administrator account is active. This is the default status.
  - **Disabled:** The administrator account is disabled.
  - **Locked:** The administrator account is locked because the password has been entered incorrectly a specified number of times.
  - **Inactive:** The administrator account is inactive because the account has not been logged in to for a specified period of time. For more information about configuring security features, see [Managing Security Operations](#).



4. Optionally, click **Next** to add extensible attributes to the admin account. For information, see [About Extensible Attributes](#).

5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Managing Passwords

Superusers can define requirements for the passwords of local admins according to your organization's policies. In addition to specifying the minimum password length, you can define rules that specify the character types that are allowed in the password. You can also specify whether passwords expire, their duration, and when reminders are sent to the users. Additionally, you can specify whether the history of used password needs to be stored, and you can require admins to change their passwords when they first log in or after their passwords are reset.

You set the requirements at the Grid level, so they apply to all local admins who log in to the Grid. You can also set the requirements at the standalone system level. The requirements that you define appear in the User Profile of all local admins and when users are required to change their password.

To define the password requirements for local admins:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab. Expand the Toolbar and select **Grid Properties -> Edit**.  
or,  
**Standalone system:** From the **System** tab, select the **System Manager** tab. Expand the Toolbar and select **System Properties Editor**.
2. In the editor, select the **Password** tab and complete the following:
  - **Minimum Password Length:** Specify the minimum number of characters that are required in a password.
  - **Password Complexity:** You can set up some requirements around how users compose a password by specifying the category and the number of characters and/or symbols the password must contain. The default is 0 for all categories, which means the password is not required to contain those characters. Specify the minimum number of characters the password must contain for the following:
    - **lowercase characters [a-z]**
    - **uppercase characters [A-Z]**
    - **numeric characters [0-9]**
    - **symbol characters.** Allowed characters are: ! @ # \$ % ^ & \* ( )
    - **character changes from previous passwords.** To discourage users from reusing previous passwords, you can require a minimum change of characters from previous passwords.
    - **password encryption.** Passwords with more than 64 characters are not encrypted.
  - **Password must expire:** Specify the number of days after which the password must expire and the number of days before which NIOS must send a reminder to the user that the password will expire.
  - **Enforce Password History:** Select this checkbox to store the history of used passwords in the NIOS database. This option is disabled by default.
    - In the **Remember last passwords** field, specify the number of passwords to be stored. You can specify a value from 1 to 20. The default value is 5.
  - **Minimum password age:** Specify the minimum number of days the password must be active before the user can attempt to change it. You can specify a value between 0 to 9998. The recommended value is 2.

Note that if the **Password must expire** checkbox is enabled, you must set the **Minimum password age** to a value less than the password expiration interval value. Superusers can override the **Minimum password age** and reset the passwords of local admins.

- **Force password change at next login:** Select this checkbox to force all new users to change their passwords when they log in for the first time, and to force existing users whose passwords were reset by superusers or whose passwords were just reset to change their passwords.



The "force password change at next login" feature does not apply to admin users in the **fireeye-group**. These users will not be prompted to change their passwords at the next login. Their original passwords continue to work. For information about FireEye integrated RPZs, see [About FireEye Integrated RPZs](#).

3. Click **Save & Close**.

## Modifying and Deleting Admin Accounts

You can modify and delete admin accounts that you create, but you can only partially modify the default superuser account "admin"—and only when you are logged in as a superuser account. Furthermore, because there must always be a superuser account on the appliance, you can only remove the default "admin" account after you create another superuser account.

To modify an admin account:

1. From the **Administration** tab, select the **Administrators** tab -> **Admins** tab -> *admin\_account* checkbox, and then click the Edit icon.  
or  
From the **Administration** tab, select the **Administrators** tab -> **Groups** tab -> *admin\_group* -> *admin\_account* checkbox, and then click the Edit icon.
2. The *Administrator* editor provides the following tabs from which you can modify data:
  - **General**: In the **General Basic** tab, modify the data of the admin account.



### Note

If the **Use AWS SSH authentication keys** option was previously disabled and is allowed when modifying an existing admin account, then password-only authentication is blocked. If the **Use AWS SSH authentication keys** option was earlier enabled and is now disabled, then password-only authentication is allowed.

Note that if the **Use AWS SSH authentication keys** option was previously disabled and is allowed when modifying an existing admin account, then password-only authentication is blocked. If the **Use AWS SSH authentication keys** option was earlier enabled and is now disabled, then password-only authentication is allowed.

On the **General Advanced** tab, complete the following:

- **Time Zone**: Select a time zone from the drop-down list if you want to specify the time zone for the administrator. By default, the appliance automatically detects the time zone from the management system that the administrator uses to connect to the appliance. The appliance uses this time zone when it displays the timestamps for relevant data.
- **Enable Certificate Authentication**: Select the checkbox to enable the certificate authentication service. You must also specify the serial number of the client certificate and associate a CA certificate that signs the client certificate. For more information, see [Enabling Certificate Authentication Service for a User](#).
- **Extensible Attributes**: Add and delete extensible attributes that are associated with the admin account. You can also modify the values of the extensible attributes. For information, see [About Extensible Attributes](#).

3. Save the configuration and click **Restart** if it appears at the top of the screen.

To delete an admin account:

1. From the **Administration** tab, select the **Administrators** tab -> **Admins** tab -> *admin\_account* checkbox, and then click the Delete icon.  
or  
From the **Administration** tab, select the **Administrators** tab -> **Groups** tab -> *admin\_group* -> *admin\_account* checkbox, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.

When you remove a Grid member from the Grid, local admin accounts are not removed and you will still be able to see these admin accounts.

## Authenticating Admin Accounts Using TACACS+

You can configure NIOS to authenticate admins against TACACS+ (Terminal Access Controller Access-Control System Plus) servers. TACACS+ provides separate authentication, authorization, and accounting services. To ensure reliable delivery, it uses TCP as its transport protocol, and to ensure confidentiality, all protocol exchanges between the

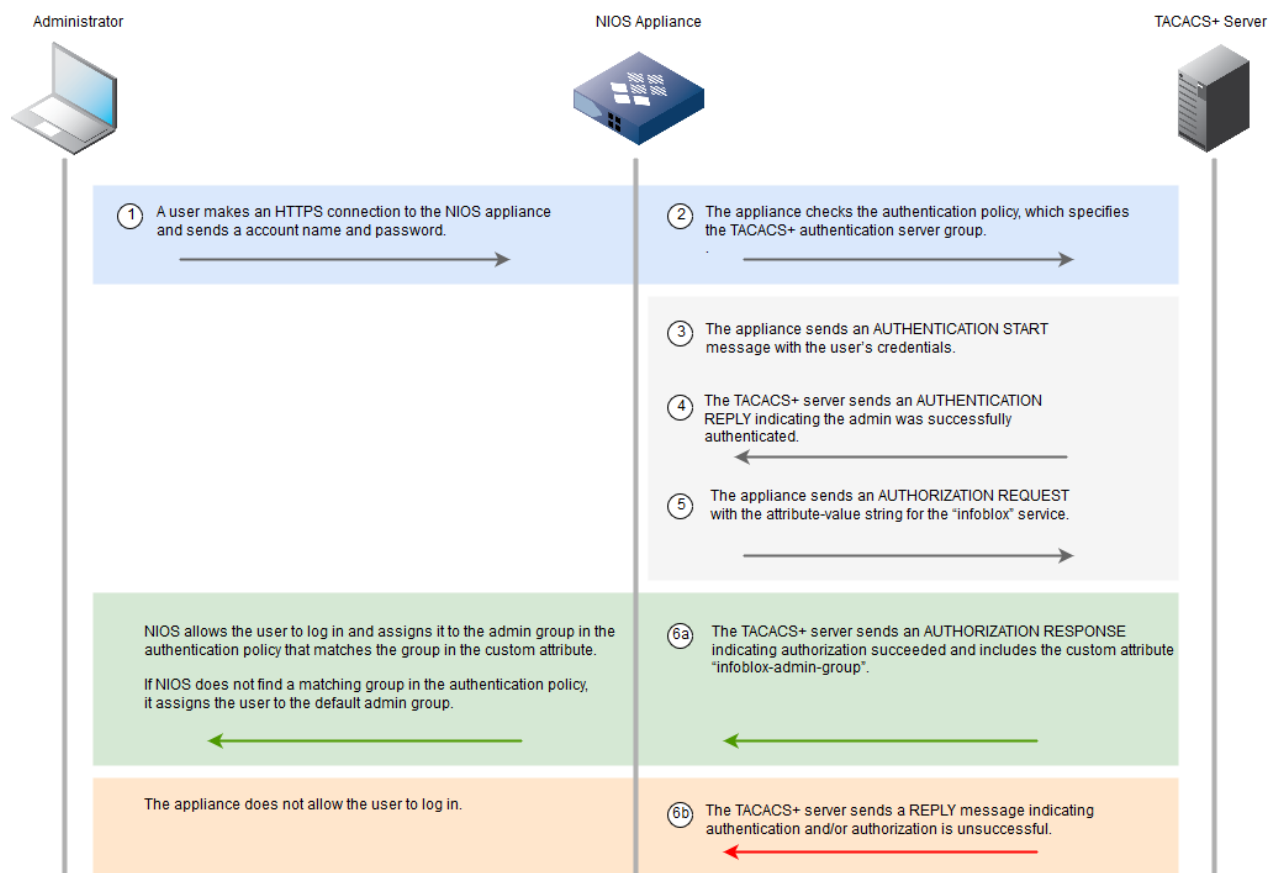
TACACS+ server and its clients are encrypted. For detailed information about TACACS+, refer to the Internet draft <https://tools.ietf.org/html/draft-grant-tacacs-02>.

In addition, you can configure a custom service, infoblox, on the TACACS+ server, and then define a user group and specify the group name in the custom attribute infoblox-admin-group. Ensure that you apply the user group to the custom service infoblox. On NIOS, you define a group with the same name and add it to the authentication policy.

Then when the TACACS+ server responds to an authentication and authorization request and includes the infoblox-admin-group attribute, NIOS can match the group name with the group in the authentication policy and automatically assign the admin to that group.

The following figure illustrates the TACACS+ authentication and authorization process when PAP/CHAP authentication is used.

### TACACS+ Authentication



### TACACS+ Accounting

When you enable TACACS+ accounting, NIOS sends the TACACS+ accounting server a TACACS+ accounting event with the same information that it sends to the Audit Log for any user command/event. NIOS sends an accounting start packet when a user first logs in successfully using TACACS+ authentication, and it sends an accounting STOP packet when a user logs out of the GUI or CLI or when a GUI or CLI session times out. If a product restarts or software failure occurs, NIOS drops any outstanding accounting packets. Note that audit log entries that are greater than 3,600 characters are truncated in accounting events sent to TACAS+ servers.

### Configuring TACACS+

Complete the following tasks to enable NIOS and the TACACS+ servers to communicate. On each TACACS+ server that you are adding to the authentication server group:

- For Windows TACACS+ servers, add the NIOS appliance as an AAA client. This step is not required for LINUX TACACS+ servers.
- Determine which user group on the TACACS+ server is used to match the admin group in NIOS, and then configure the following settings for the user group:
  - Add "infoblox" as a custom service.
  - Define the custom attribute for the group, in the format: **infoblox-admin-group= *group\_name***. For example, **infoblox-admin-group=remoteadmins1**. The group name can have a maximum of 64 characters.

On the NIOS appliance:

- Create a TACACS+ authentication server group. You can create only one TACACS+ server group. For more information, see [Configuring a TACACS+ Authentication Server Group](#) below.
- Create the local admin group in NIOS that matches the user group on the TACACS+ server. Note that the NIOS admin group name must match the group name specified in the TACACS+ server and in the custom attribute. For example, if the custom attribute is **infoblox-admin-group=remoteadmins1**, then the admin group name must be **remoteadmins1**. In addition, you can designate a default admin group for remote admins. For information about configuring group permissions and privileges, see [About Admin Groups](#).
- In the authentication policy, add the newly configured TACACS+ server group and the TACACS+ admin group name. See [Defining the Authentication Policy](#) for more information about configuring an admin policy.

## Configuring a TACACS+ Authentication Server Group

You can add multiple TACACS+ servers to the TACACS+ authentication server group. NIOS sends authentication requests to the TACACS+ servers in the order they are listed. NIOS sends authentication requests to the first server on the list. If that server is unreachable or generates an error, then NIOS sends the request to the next server in the list that has not been previously queried, and so on. NIOS logs an error message in syslog if all servers have been queried and they all generate errors or are unreachable.

To configure a TACACS+ authentication server group:

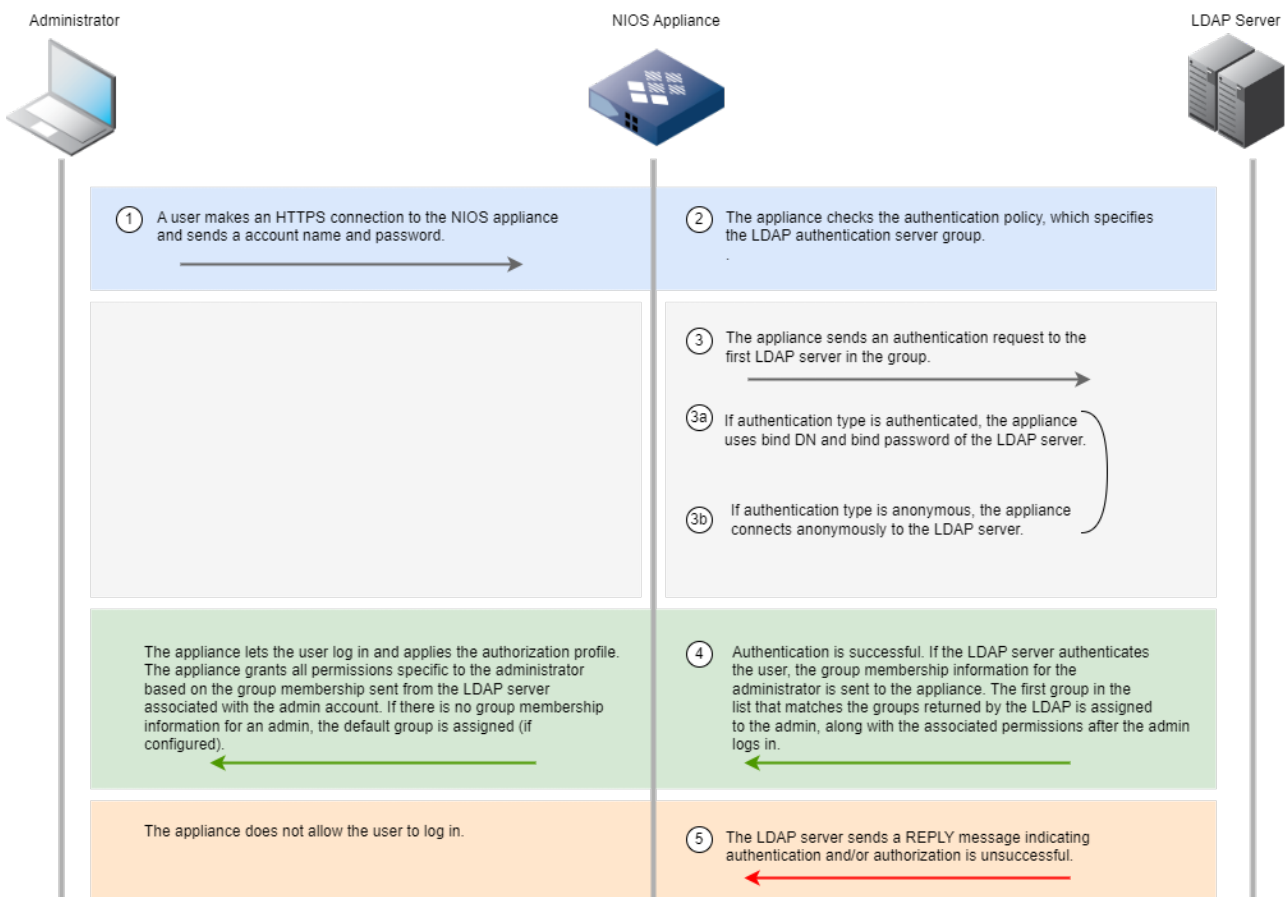
1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the **TACACS+ Services** subtab and click the Add icon.
3. In the *Add TACACS+ Service* wizard, complete the following:
  - **Name:** Enter a name for the server group.
  - **TACACS+ Servers:** Click the Add icon and complete the following:
    - **Server Name or IP address:** The name or IP address of the TACACS+ server.
    - **Comment:** You can enter additional information about the server.
    - **Port:** The TCP destination port for TACACS+ communication. This port is used for authentication, accounting and authorization packets. The default is port 49.
    - **Authentication Type:** Select **ASCII**, **PAP** or **CHAP**. The default is **CHAP**.
    - **Shared Secret:** The shared key that the NIOS appliance and the TACACS+ server use to encrypt and decrypt messages.
    - **Enable Accounting:** Select this to enable NIOS to send accounting information to the TACACS+ server.
    - **Connect through Management Interface:** Select this checkbox to enable the appliance to use the MGMT port to communicate with the TACACS+ server. Ensure that the MGMT port is configured. Otherwise, the appliance will use the LAN interface
    - **Disable Server:** Select this to prevent queries from being sent to this server. You can retain the configuration, but disable the service.  
Click **Test** to test the configuration. Click **Add** to add the TACACS+ server to the list.  
When you add multiple TACACS+ servers, the appliance lists the servers in the order you added them. This list also determines the order in which the NIOS appliance attempts to contact a TACACS+ server. You can move a server up or down the list by selecting it and clicking the up or down arrow.
    - **Authentication/Authorization:** Optionally, modify the authentication and authorization settings. These settings apply to all TACACS+ servers that you configure on the NIOS appliance.
      - **Timeout(s):** Specify the number of seconds or milliseconds that the appliance waits for a response from the TACACS+ server before it tries to contact it again. The amount of time before the server is retried. The default and minimum is 5000, and the maximum is 60000.

- **Retries:** Specify how many times NIOS attempts to contact a TACACS+ server and fails before it tries to contact the next server on the list. The default is 0. The maximum is 5.
  - **Accounting:** Optionally, modify the Accounting settings.
    - **Timeout(s):** Specify the number of seconds or milliseconds that the appliance waits for a response from the TACACS+ server. The amount of time before the server is retried. The default and minimum is 1000, and the maximum is 30000.
    - **Retries:** Specify how many times the appliance attempts to contact an accounting TACACS+ server and fails before it tries to contact the next accounting server on the list. The default is 0. The maximum is 5.
  - **Comment:** Enter additional information about the service.
  - **Disable:** Select this to retain an inactive TACACS+ authentication service profile.
4. Save the configuration.

## Authenticating Admins Using LDAP

LDAP (Lightweight Directory Access Protocol) is an internet protocol for accessing distributed directory services. The NIOS appliance can authenticate admin accounts by verifying user names and passwords against LDAP. The NIOS appliance queries the LDAP server for the group membership information of the admin. The appliance matches the group names from the LDAP server with the admin groups in its local database. It then authorizes services and grants the admin privileges, based upon the matching admin group on the appliance. The following figure illustrates the LDAP authentication process.

### Authenticating using an LDAP server



### Authentication Protocols

When you configure the NIOS appliance to authenticate admins against an LDAP server group, you must specify the authentication protocol of each LDAP server, which can be either anonymous or authenticated. The NIOS appliance connects anonymously to the LDAP server when the authentication type is anonymous. With authenticated type, the

NIOS appliance connects using the bind DN and bind password defined for that server.

You can configure one of the following modes to send the authentication request to the LDAP server:

- **Ordered:** In this mode, the authentication request is sent to the first server in the list. The authentication request is sent to the next server only when the first server is out of service or unavailable.
- **Round Robin:** In this mode, the first authentication request is sent to a server chosen randomly in a group. If there is no response from the server, continued attempts are performed sequentially until it selects the last server in the list. Then it starts with the first server in the list and continues the selection process until all the servers have been attempted.

You can also specify the authentication type, for admins who belong to specific groups. The NIOS appliance uses the selected group authentication type to query the LDAP server and retrieve the group names to which the admin belongs. In LDAP, you can group users by any custom object classes. Example: `objectclass groupofNames`, `objectclass posixGroup`, etc. In NIOS, when you select Member Group Attribute as the group authentication type, the appliance uses custom LDAP group attributes to query the LDAP server and retrieve the group names for authentication. Example: `memberOf`, `isMemberOf`, etc. When you select Posix Group as the authentication type, the appliance uses "memberuid" and "objectClass" to query the server and retrieve the group names for authentication.

## Configuring LDAP

Do the following to configure NIOS to use one or more LDAP server groups to authenticate administrators:

- Configure at least one LDAP authentication server group. For more information, see [Configuring an LDAP Server Group](#) below.
- Define admin groups for the admins that are authenticated by the LDAP servers and specify their privileges and settings. The group names in NIOS must match the admin group names on the LDAP server. For more information about defining admin groups, see [About Admin Groups](#).
- In the authentication policy, add the LDAP server groups and the admin groups that match those on the LDAP server. You can also designate an admin group as the default group for remote admins. NIOS assigns admins to this group when it does not find a matching group for a remote admin. For more information about configuring the policy, see [Defining the Authentication Policy](#).

## Configuring an LDAP Server Group

You can add one or more LDAP servers to an LDAP group for redundancy. The NIOS appliance tries to connect with the LDAP server based on the method you configure for the authentication request. If it does not receive a response within the configured timeout period and has tried to connect the specified retry value, then it tries the next LDAP server on the list. The appliance makes a syslog entry when it fails to connect to any of the servers in the group and sends an SNMP trap and an email notification if configured.

To configure an LDAP server group on the NIOS appliance:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the Add icon in the **LDAP Services** subtab.
3. In the *Add LDAP Authentication Service* wizard, complete the following:
  - **Name:** Enter the name of the server group.
  - **LDAP Servers:** Click the Add icon and enter the following:
    - **Server Name or IP Address:** Enter the FQDN (fully-qualified domain name) of the server or enter the IPv4/IPv6 address.
    - **LDAP Version:** Select the LDAP version. The NIOS appliance supports both LDAPv2 and LDAPv3. The default LDAP version is v3.
    - **Base DN:** Enter the base DN (Distinguished Name) value. All entries stored in an LDAP directory have a unique DN.
    - **Authentication Type:** Select the authentication type from the drop-down list. The supported authenticated types are as follows:
      - **Anonymous:** Select this to connect to the LDAP server anonymously. This is selected by default.
      - **Authenticated:** Select this to connect using the bind DN and bind password defined for that server.
        - **Bind User DN:** Enter the bind user DN.

- **Bind Password:** Enter the bind password.
- **Encryption:** Select the encryption type from the drop-down list.
  - **SSL:** This is selected by default. All the network traffic is encrypted through an SSL (Secure Sockets Layer) protocol. The appliance automatically updates the authentication port to 636 for SSL. You must upload a CA certificate that verifies the LDAP server certificate. Click **CA Certificates** to upload the certificate. In the *CACertificates* dialog box, click the Add icon, and then navigate to the certificate to upload it.
  - **NONE:** Select this to unencrypt the connection. Note that Infoblox strongly recommends that you select the SSL option to ensure the security of all communications between the server and the member.
- **Network Port:** Enter the authentication port number on the LDAP server to which the appliance sends authentication requests. The default value is 636. When you select NONE from the Encryption drop-down list, the appliance automatically updates the authentication port to 389.
- **Comment:** Enter useful information about the LDAP server.
- **Connect through Management Interface:** Select this so that the NIOS appliance uses the MGMT port for administrator authentication communications with just this LDAP server.
- **Disable Server:** Select this to disable the LDAP server if, for example, the connection to the server is down and you want to stop the NIOS appliance from trying to connect to this server. You cannot disable the only server in a group if it is already being used by the remote authentication policy.
- Click **Test** to test the configuration. If the NIOS appliance connects to the LDAP server using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the server, the appliance displays a message indicating an error in the configuration.
- Click **Add** to add the LDAP server to the group.
 

When you add multiple LDAP servers, the appliance lists the servers in the order you added them. This list also determines the order in which the NIOS appliance attempts to contact an LDAP server. You can move a server up or down the list by selecting it and clicking the up or down arrow.

You can also delete a server by selecting it and clicking the Delete icon.
- **Server Timeout(s):** Specify the number of seconds that the appliance waits for a response from the LDAP server. The default value is 5 seconds.
- **Retries:** Specify how many times the appliance attempts to contact an authentication LDAP server. The default value is 5.
 

If you have configured multiple LDAP servers for authentication and the NIOS appliance fails to contact the first server in the list, it tries to contact the next server after completing the specified number of attempts, and so on.
- **Mode:** Specifies the order in which a Grid member connects to an LDAP server.
  - **Ordered List:** The Grid member always selects the first LDAP server in the list when it sends an authentication request. It queries the next server only when the first server is considered down. This is the default.
  - **Round Robin:** The Grid member sends the first authentication request to a server chosen randomly in a group. If there is no response from the server, the Grid member selects the next server in the group. Continued attempts are performed sequentially until it selects the last server in the group. Then it starts with the first server in the group and continues the selection process until all the servers have been attempted.
- **Recovery Interval:** Specify the number of seconds that the appliance waits to recover from the last failed attempt in connecting to an LDAP server. Select the time unit from the drop-down list. The default is 30 seconds. This is the time interval that NIOS waits before it tries to contact the server again since the last attempt when the appliance could not connect to the LDAP server or when the LDAP server did not send a reply within the configured response timeouts and retry attempts.
- **Group Authentication Type:** Select the group authentication type for LDAP authentication service from the drop-down list. By default, **Member Group Attribute** authentication type is selected. When you select **Member Group Attribute**, you can specify custom LDAP group attribute in the **Group Membership Attribute** field. For example, memberOf, isMemberOf, etc. The appliance uses this attribute to retrieve the group names to which the admin belongs. When you select **Posix Group**, the appliance uses "memberuid" and "objectClass" to retrieve the group names to which the admin belongs.
- **Group Membership Attribute:** Specify the LDAP group attribute (such as "memberOf" and "isMemberOf"). This is used to query the server and retrieve the group names to which the admin belongs. This field is enabled only when you select Member Group Attribute in the **Group Authentication Type** drop-down list. The default value is **memberOf**.



- **LDAP Search Scope:** To search for an admin user name in the LDAP directory, select one of the following LDAP search scope:
    - **Base:** Specify Base to perform search only on base in the LDAP directory. This is the top level of the LDAP directory tree.
    - **One Level:** Specify One Level to perform search on base DN and one level below the base in the LDAP directory.
    - **Subtree:** Specify Sub tree to perform search on base and all the entries below the base DN in the LDAP directory.  
The default value is One Level.
  - **User ID:** Specify the attribute associated with the user ID object in the LDAP server, such as "uid" and "cn". This attribute is used to match the NIOS user name.
  - **Map LDAP Field to Extensible Attribute (for Captive Portal Users only):** If you configure the LDAP authentication server group to authenticate the captive portal users, you can map an LDAP attribute value to an existing extensible attribute. This mapping is optional. By doing so, the LDAP attribute value will be queried from the LDAP server once the captive portal user authentication is successful. The attribute value received from the LDAP server is mapped to the corresponding extensible attribute. NIOS updates or creates a MAC address filter depending on the captive portal user or the client's hardware and name. Click the Add icon and enter the following:
    - **LDAP Field:** Enter the LDAP attribute. This attribute is queried in the LDAP directory server.
    - **Extensible Attributes:** Select an attribute from the drop-down list. The drop-down list displays only the extensible attributes configured with attribute type as string. Infoblox recommends that you avoid confidential data while mapping extensible attribute to an LDAP attribute because this data is visible in the extensible attribute field of the corresponding MAC address filter.  
Note that mapping an extensible attribute to an LDAP attribute must be unique for a given LDAP server. Attribute not defined in the LDAP directory for a given user is considered as null and is mapped to the corresponding extensible attribute with a default value. The default value of extensible attribute is Not Found. This default value is not configurable and they do not cause the authentication to fail.
  - **Comment:** Enter useful information about the LDAP server group.
  - **Disable:** Select this to disable the LDAP authentication server group. Note that you cannot disable an LDAP group if it is already being used to authenticate one or more administrators and/or captive portal users.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## About Admin Accounts

A user must have an admin account to log in to the NIOS appliance. Each admin account belongs to an admin group, which contains roles and permissions that determine the tasks a user can perform. For information, see [About Admin Groups](#).

When an admin connects to the appliance and logs in with a username and password, the appliance starts a two-step process that includes both authentication and authorization. First, the appliance tries to authenticate the admin using the username and password. Second, it determines the authorized privileges of the admin by identifying the group to which the admin belongs. It grants access to the admin only when it successfully completes this process.

The NIOS appliance can authenticate users that are stored on its local database as well as users stored remotely on an Active Directory domain controller, a RADIUS server, a TACACS+ server or an LDAP server. The group from which the admin receives privileges and properties is stored locally.

NIOS can authenticate users based on X.509 client certificates irrespective of the client certificate source. For example, smart card holders such as U.S. Department of Defense CAC users and PIV card holders. The status of these certificates is stored remotely on OCSP (Online Certificate Status Protocol) responders. NIOS uses two-factor authentication to validate these users. For more information about two-factor authentication and how to configure it, see [Authenticating Admins Using Two-Factor Authentication](#).

The tasks involved in configuring administrator accounts locally and remotely are listed in *Storing Admin Accounts Locally and Remotely* table



NIOS Appliance		RADIUS server/AD Domain Controller/TACAS+ server/ LDAP server/Certificate authentication service
<b>To store admin accounts locally</b>	<ul style="list-style-type: none"> <li>• Use the default admin group ("admin-group") or define a new group</li> <li>• Set the privileges and properties for the group</li> <li>• Add admin accounts to the group</li> </ul>	
<b>To store admin accounts remotely</b>	<ul style="list-style-type: none"> <li>• Configure communication settings with a RADIUS server, an Active Directory domain controller, TACACS+ server, or LDAP server</li> </ul> <p>If you use admin groups on the RADIUS server, Active Directory domain controller, TACACS+ server, or LDAP server:</p> <ul style="list-style-type: none"> <li>• Configure admin groups that match the remote admin groups</li> <li>• Set the privileges and properties for the groups</li> </ul> <p>If you do not use admin groups on the RADIUS server, Active Directory domain controller, TACACS+ server, or LDAP server:</p> <ul style="list-style-type: none"> <li>• Assign an admin group as the default</li> </ul>	<ul style="list-style-type: none"> <li>• Configure communication settings with the NIOS appliance</li> </ul> <p>If you use admin groups:</p> <ul style="list-style-type: none"> <li>• Import Infoblox VSAs (vendor-specific attributes) (if RADIUS)</li> <li>• Define an admin group with the same name as that on the NIOS appliance</li> <li>• Define admin accounts and link them to an admin group</li> </ul> <p>If you do not use admin groups:</p> <ul style="list-style-type: none"> <li>• Define admin accounts</li> </ul>

The admin policy defines how the appliance authenticates the admin: with the local database, RADIUS, Active Directory, TACACS+, or LDAP. You must add RADIUS, Active Directory, TACACS+, or LDAP as one of the authentication methods in the admin policy to enable that authentication method for admins. See [Defining the Authentication Policy](#) for more information about configuring the admin policy.

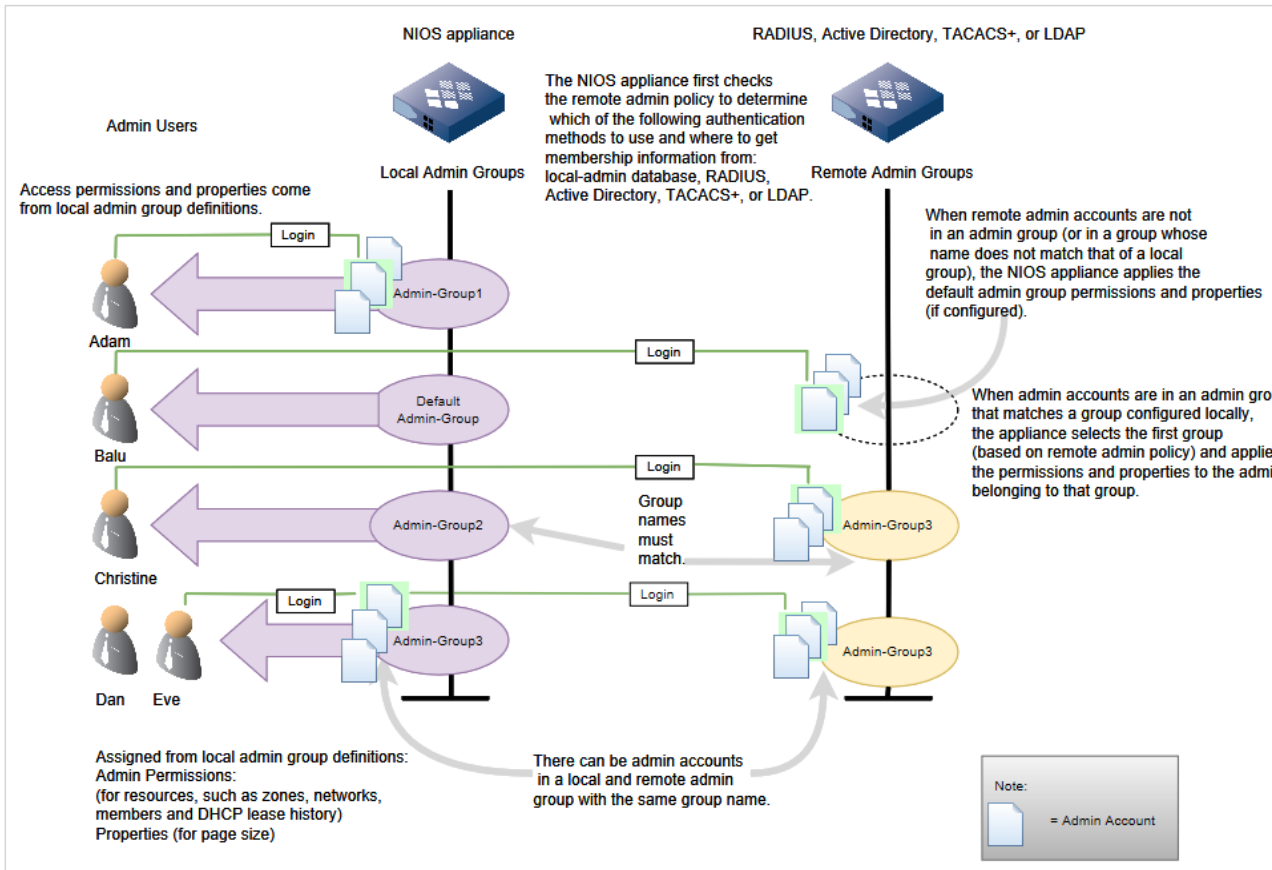


**Note**

Local passwords are stored in the database as part of the user object. Values for passwords are stored after applying a random salt and hashed with SHA-128.

The following figure illustrates the relationship of local and remote admin accounts, admin policy, admin groups, and permissions and properties.

*Privileges and Properties Applied to Local and Remote Admin Accounts*



Complete the following tasks to create an admin account:

1. Use the default admin group or create an admin group. See [About Admin Groups](#).
2. Define the administrative permissions of the admin group. See [About Administrative Permissions](#).
3. Create the admin account and assign it to the admin group.
  - To add the admin account to the local database, see [Creating Local Admins](#).
  - To configure the appliance to authenticate the admin account stored remotely, see [About Remote Admins](#).

## Defining the Authentication Policy

The authentication policy defines which authentication server groups the appliance uses to authenticate admins and lists the local admin groups that map to the remote admin groups.

By default, the appliance provides the "Local Admin" service for authenticating users against the local database. You cannot modify or delete this default service.

## Configuring a List of Authentication Server Groups

To enable NIOS to use multiple authentication server groups, define a prioritized list as follows:

1. From the **Administration** tab, select the **Administrators** tab -> **Authentication Policy** tab.
2. From the **Authenticate users against these services in this order** section, click the Add icon to add an authentication server group.
3. Select one of the following in the *Add Authentication Service* section:
  - **Active Directory**: Select this to add an AD authentication server group, and then select a group from the drop-down list.

- **RADIUS:** Select this to add a RADIUS authentication server group, and then select a group from the drop-down list.
- **TACACS+:** Select this to add the TACACS+ authentication server group, and then select a group from the drop-down list.
- **LDAP:** Select this to add the LDAP authentication server group, and then select a group from the drop-down list.
- **Certificate Authentication Service:** Select this to add a certificate authentication service, and then select a service from the drop-down list.
- **SAML:** Select this to use SAML SSO as the authentication service.

4. Click **Add**.

You can reorder the list by selecting an authentication server group and moving it up or down the list using the arrow keys.



#### Note

Using a remote authentication service causes high memory utilization on discovery members. However, this does not affect the operation of other processes.

## Configuring a List of Remote Admin Groups

In order for NIOS to assign a remote admin to the correct group, you must list the admin groups in the local database that match the remote admin groups. You can also define a default admin group to which NIOS assigns remote users with no admin groups listed.

The appliance matches a remote admin to a group in the order the groups are listed. When the appliance receives information that an admin belongs to one or more groups, the appliance assigns the user to the first group in the list that matches. It assigns the admin to the default group, if specified, if no groups are returned by the authentication server, or if the appliance does not find a group in the local database that matches the group returned by the authentication server.

To configure the remote admin group list:

1. From the **Administration** tab, select the **Administrators** tab -> **Authentication Policy** tab.
2. In the **Authentication Server Groups is the authority for** section, select one of the following:
  - **Remote users, their passwords and their groups ownership:** Select this to use the authentication server groups to define the list of remote users, their passwords and their group ownerships. This is selected by default. You can add the list of admin groups to map the remote admin group to a local group in the **Map the remote admin group to the local group in this order** section.
  - **Passwords of Local users:** Select this to authenticate a remote user when you do not know to which remote admin group the user belongs.
  - **Map the remote admin group to the local group in this order:** In order for the appliance to assign a remote admin to the correct group, you must list the admin groups in the local database that match the remote admin groups. The appliance matches a remote admin to a group in the order the groups are listed. You can also define a default admin group to which NIOS assigns remote users with no admin groups listed. This section is disabled when you select **Passwords of Local users** in the **Authentication Server Groups is the authority for** section.

When the appliance receives information that the admin belongs to one or more groups, the appliance selects the first group in the list that matches, and assigns that group to the admin. If no groups are returned by the authentication server, the default group is assigned (if specified).

Complete the following to configure the remote admin group list:

- Click the Add icon to add an admin group to the list. In the *Admin Group Selector* dialog box, select an admin group. Use Shift+click and Ctrl+click to select multiple admin groups.
- You can reorder the list by selecting an admin group and using the arrow keys to move it up or down the list.
- **Assign user to this group if remote admin group cannot be found:** Click **Select** to assign a user to a specific admin group if the remote admin group is not found. In the *Admin Group Selector* dialog

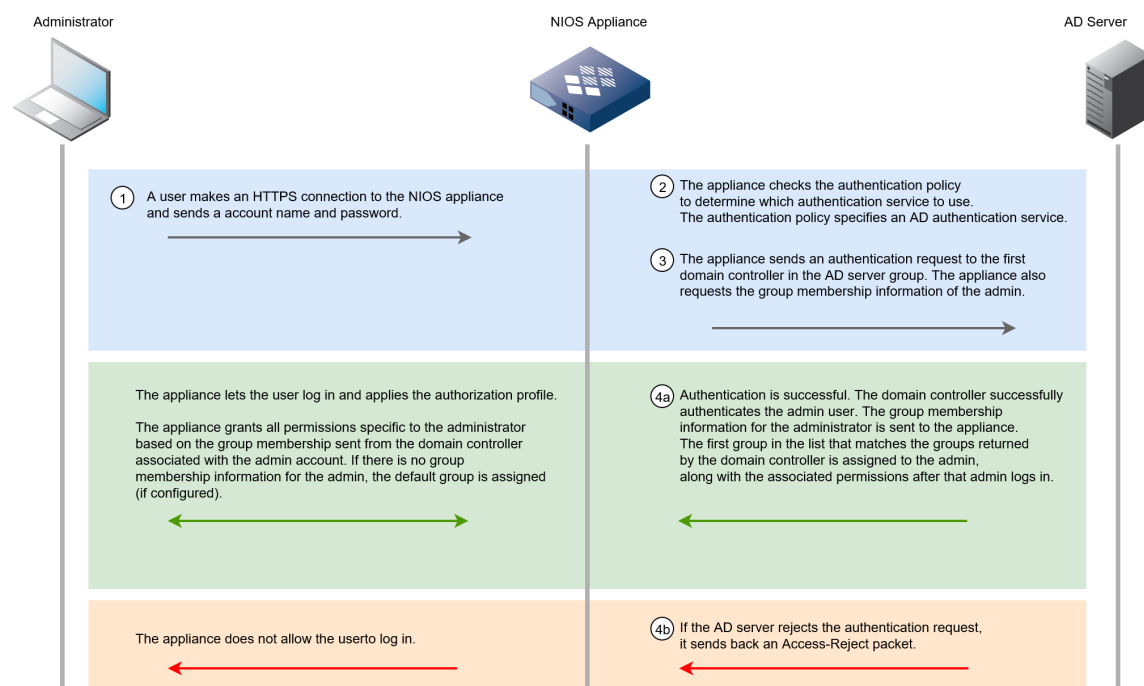
box, select an admin group. You can also click **Clear Selection** to clear the displayed member and select a new one.

## Authenticating Admins Using Active Directory

Active Directory™ (AD) is a distributed directory service that is a repository for user information. The NIOS appliance can authenticate admin accounts by verifying user names and passwords against Active Directory. In addition, the NIOS appliance queries the AD domain controller for the group membership information of the admin. The appliance matches the group names from the domain controller with the admin groups on its local database. It then authorizes services and grants the admin privileges, based upon the matching admin group on the appliance.

The following figure illustrates the Active Directory authentication process.

### Authentication Using a Domain Controller



To configure NIOS to authenticate administrators using Active Directory domain controller groups, you must first configure user accounts on the domain controller.



#### Note

Do not create Microsoft user accounts with the following characters: "", +, ,, ;, <, =, >, \. Microsoft does not allow creating users with these characters and such characters will be replaced by an underscore \_.

Then, on the NIOS appliance, do the following:

- Configure one or more AD authentication server group on the appliance and add AD domain controllers to the group. For information about configuring an AD authentication service group for admins, see [Configuring an Active Directory Authentication Service Group](#).
- If you configured admin groups on the AD controller, you must create those same groups on the NIOS appliance and specify their privileges and settings. Note that the admin group names must match those on the AD domain controller. You can specify a default group as well. The NIOS appliance assigns admins to the default group if none of the admin groups on the NIOS appliance match the admin groups on the AD domain controller or if there are no other admin groups configured. For information about configuring group permissions and privileges, see [About Admin Groups](#).

- Add the newly configured Active Directory service to the list of authentication services in the admin policy, and add the admin group names as well. See [Defining the Authentication Policy](#) for more information about configuring an admin policy.

## Configuring an Active Directory Authentication Service Group

You can add multiple domain controllers to an AD authentication server group for redundancy. The NIOS appliance tries to connect with the first domain controller on the list. If it is unable to connect, it tries the next domain controller on the list, and so on.

To configure an Active Directory authentication server group on the NIOS appliance:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the **Active Directory Services** subtab and click the Add icon.
3. In the *Add Active Directory Authentication Service* wizard, complete the following:
  - **Name:** Enter a name for the service.
  - **Active Directory Domain:** Enter the AD domain name.
  - **Domain Controllers:** Click the Add icon and complete the following to add an AD domain controller:
    - **Server Name or IP Address:** Enter the FQDN or the IP address of the AD server that is used for authentication.
    - **Comment:** Enter additional information about the AD server.
    - **Authentication Port:** Enter the port number on the domain controller to which the appliance sends authentication requests. The default is 389.
    - **Encryption:** Select **SSL** from the drop-down list to transmit through an SSL (Secure Sockets Layer) tunnel. When you select SSL, the appliance automatically updates the authentication port to 636. Infoblox strongly recommends that you select this option to ensure the security of all communications between the NIOS appliance and the AD server. If you select this option, you must upload a CA certificate from the AD server. Click **CA Certificates** to upload the certificate. In the *CA Certificates* dialog box, click the Add icon, and then navigate to the certificate to upload it.
    - **Connect through Management Interface:** Select this so that the NIOS appliance uses the MGMT port for administrator authentication communications with just this AD server.
    - **Disable server:** Select this to disable an AD server if, for example, the connection to the server is down and you want to stop the NIOS appliance from trying to connect to this server.
    - Click **Test** to test the configuration. If the NIOS appliance connects to the domain controller using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the server, the appliance displays a message indicating an error in the configuration.
    - Click **Add** to add the domain controller to the group.

When you add multiple domain controllers, the appliance lists the servers in the order you added them. This list also determines the order in which the NIOS appliance attempts to contact a domain controller. You can move a server up or down the list by selecting it and clicking the up or down arrow.

You can also delete a domain controller by selecting it and clicking the Delete icon.

- **Timeout(s):** The number of seconds that the NIOS appliance waits for a response from the specified authentication server. The default is 5.
  - **Comment:** Enter additional information about the service.
  - **Disable:** Select this to retain an inactive AD authentication service profile.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Enabling Active Directory Authentication for Nested Groups

Windows servers support nesting groups in which you can add a group of admin users as a member of another group. Nested groups consolidate admin accounts and help reduce the number of permissions required for individual users or groups. In NIOS, you can enable a nested group query so the appliance can recursively look up and use the AD authentication service to authenticate members or admin accounts. These members or admin accounts can be part of the default nested group, outside of the default nested group, or located within a non-default custom organizational unit.

When an admin belongs to multiple paths of hierarchy, you can enable nested group query in order to apply the AD authentication service hierarchically in a parent-child structure. This enables the NIOS appliance to apply the AD

authentication service to all the groups of which an admin is a member. For example, if User 1 is a member of the default nested Group C, and Group C is a member of Group B, and Group B is a member of Group A, then the authentication service is applicable to all the groups of which User 1 is a member. In this example, the appliance performs a recursive lookup in Group C, Group B, and Group A while authenticating User 1.

You can also define multiple organizational units and add non-default AD admins and groups to these units.



#### Note

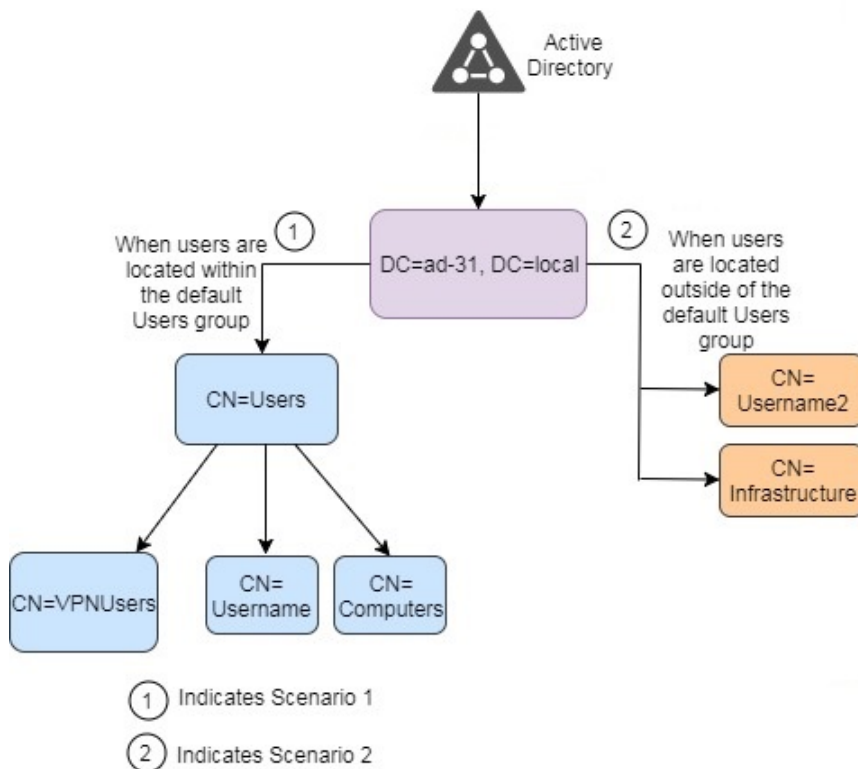
- Microsoft recommends that you create all non-default users and groups in different organizational units to apply Group Policy Objects and prevent corruption or misuse of critical default accounts and groups.
- In Active Directory, objects are referred to by the DN (Distinguished Name), which contains CN (Common Name), OU (Organizational Unit), and DC (Domain Component) that are delimited by commas and indicate where the object resides in an AD hierarchy.

Infoblox supports AD authentication for nested groups in the following scenarios:

- **Scenario 1:** When the user is located within the default **Users** group.  
In this example, the DC ad-31.local contains CN=Users that is the default user group or the container. Users Username, VPNUUsers, and Computers are located within the default user group. When you enable a nested group query, the appliance uses the AD authentication service to authenticate the admin accounts that are within the default nested group CN=Users. The DN for users within the default user group are as follows:
  - DN: CN=Username, CN=Users, DC=ad-31, DC=local
  - DN: CN=VPNUUsers, CN=Users, DC=ad-31, DC=local
  - DN: CN=Computers, CN=Users, DC=ad-31, DC=local
- **Scenario 2:** When the user is located outside the default **Users** group.  
In this example, users CN=Username2 and CN=Infrastructure are not located within the default user group or the container. The DN for users within the DC ad-31.local are as follows:
  - DN: CN=Username2, CN=Users, DC=ad-31, DC=local
  - DN: CN=Infrastructure, CN=Users, DC=ad-31, DC=local

The following figure illustrates how you can configure default and non-default nested groups and add users to these groups in a Windows Active Directory. It contains a DC ad-31.local that can contain default users group, organizational units or individual users. When you select the **Disable Default Search Path** checkbox, the AD authentication service authenticates the admin account that is mentioned in the additional search path for a non-default organizational unit.

*Active Directory Authentication for Nested Groups*



To enable AD authentication for nested groups on the NIOS appliance:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the **Active Directory Services** subtab and click the Add icon.
3. In the *Add Active Directory Authentication Service* wizard, complete the following:
  - Nested Group Query:** This checkbox is not selected by default, meaning the nested group query is disabled. When nested group query is disabled, AD authentication service is applied to only one group of which the AD admin is a member. When you select this checkbox, AD authentication service is applied to all the nested groups of which an AD admin is a member. This setting is applicable to all the AD servers configured for the Active Directory authentication service.
4. Save the configuration.

## Changing Password Length Requirements

Password length requirements control how long a password must be for a NIOS appliance admin account. Increasing this value reduces the likelihood of hackers gaining unauthorized access.

To change password length requirements:

1. From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties* editor, select the **Password** tab.
3. Enter a number from 4 to 64 in the **Minimum Password Length** field.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Notifying Administrators

You can notify individual administrators about system status through email, or notify a group of people using an alias email address. If you have configured DNS resolution on your network, the **E-mail relay configuration** function is not



required. If you did not configure the settings on the **DNS Resolver** section, you must enter a static IP address of the target system in the **Relay Name/IP Address** field. The appliance sends e-mail to administrators when certain events occur. This functionality supports both IPv4 and IPv6 networks. Here is a list of events that trigger e-mail notifications:

- Changes to link status on ports and online/offline replication status
- Events that generate traps, except for upgrade failures (ibUpgradeFailure). For a list of events and Infoblox MIBs, see [SNMP MIB Hierarchy](#).

The appliance attempts to send the email notification once after an event. It does not try to send the notification again, if the first attempt fails. Infoblox recommends that you use the **Test Email settings** button to test the email settings and to verify that the recipient received the notification.

You can define the email settings at the Grid and member levels.

## Grid Level

To notify an administrator of an independent appliance or a Grid:

1. From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties* editor, select the **Email** tab, and then complete the following:
  - **Enable Email notification:** Select this for NIOS to send email notifications.
    - **From Email Address:** Enter the email address from which to send email notifications. If this field is blank, the member host name is used by default. For example, `no-reply@member_host_name`
    - **To Email Address:** Enter the email address of the administrator. Use an email alias to notify multiple people.
  - **Use SMTP Relay:** Select this if the NIOS appliance must send email to an intermediary SMTP (Simple Mail Transfer Protocol) server that relays it to the SMTP server responsible for the domain name specified in the email address. Some SMTP servers only accept email from certain other SMTP servers and might not allow email from the NIOS appliance. In this case, specify the DNS name or IP address of a different SMTP server that does accept email from the NIOS appliance and that will then relay it to the SMTP server that can deliver it to its destination.
    - Clear this if it is unnecessary to use an email relay server.
      - **SMTP Relay Name or Address:** If you have configured DNS resolution, enter the DNS name of the relay server.  
If DNS resolution is not configured, enter the IP address of the relay server.
3. Optionally, click **Test Email settings** to confirm this feature is operating properly.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Member Level

To define email settings for a member:

1. From the **Grid** tab, select the **Grid Manager** tab -> *member* checkbox, and then select the Edit icon.
2. In the *Grid Member Properties* editor, select the **Email** tab, and then click **Override** to override Grid-level settings.
3. Complete the email configuration as described in Grid Level above.

## Administrative Permissions for Common Tasks

Permissions for Common Tasks table lists some of the common tasks admins can perform and their required permissions. [Permission for Network Discovery](#) table lists tasks related to device discovery, with their required permissions.

All the permission tables in this chapter use the following definitions:

**RW** = Read/Write permission

**RO** = Read-only permission

### *Permissions for Common Tasks*



Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank A/AA record
For Grid and Members																				
Restart services for the entire Grid	RO																			

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank A/AA Record
Configure Grid DNS properties, configure member DNS properties, assign members to DNS objects, and restart DNS service on members		R W																		

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank AA Record
Configure Grid DHCP properties, configure member DHCP properties, assign members to DHCP objects, and restart DHCP service on members			RW																	

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank / AA record
Configure a Grid member				R W																
Restart services on a Grid member				R W																

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank A/AA record
Configure member DNS properties, assign member to DNS objects, and restart DNS service on member					R W															

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank A/AA record
Configure member DHCP properties, assign member to DHCP objects, and restart DHCP service on member					R W															
Restart member DNS service							R W													

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank AAA Record
Restart member DHCP service								R W												
Initiate and control network discovery on all networks														R W						R W
Scheduling tasks for all supported objects									R W				R W							R W

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank A/AA record
For DNS resources																				
Create, modify, and delete DNS views									R W											
View and search for DNS views									R O											



Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank A/AA Record
Create, modify, and delete DNS zones with as signed members	RW									RW										
View and search for DNS zones with as signed members	RO									RO										

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank A/AA Record
Create, modify, and delete all resource records	R W											R W								
Create and modify blank A/AA records and shared A/AA records.																				R W

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank A/AA Record
View and search for all resource records	RO											RO								
Assign member to DNS objects							RW													
For DHCP Resources																				

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank A/AA Record
Create, modify, and delete network views and their associated DNS views	RW								RW				RW							
View network properties and statistics	RO												RO							

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank A/AA Record
Create, modify, and delete networks with assigned members				R W										R W						
Create, modify, and delete networks without assigned members															R W					

Tasks	All Grid Members	Grid DNS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank A/AA record
Create, modify, and delete DHCP ranges in a specific network with assigned members				RW											RW	RW				

Tasks	All Grid Members	Grid NS Properties	Grid DHCP Properties	Specific Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	All DNS Views	All DNS Zones	All Shared Record Groups	All Resource Records	All Network Views	All Networks	Specific Network(s)	DHCP Range(s)	Fixed Addresses	Scheduling Task	Network Discovery	Adding a blank AAA Record
Create, modify, and delete fixed addresses in a specific network without assigned members															R W		R W			
Assign member to DHCP objects								R W												

### Administrative Permission for the Grid

By default, the Grid Master denies access to Grid members when a limited-access admin group does not have defined permissions. You can grant an admin group read-only or read/write permission, or deny access to all Grid members or you can grant permission to specific Grid members, as described in [Applying Permissions and Managing Overlaps](#), see [About Administrative Permissions](#).



### Note

Only superusers can modify DNS and DHCP Grid properties.

The following sections describe the types of permissions that you can set with Grid permissions:

- Administrative Permissions for Grid Members
- Administrative Permissions for Network Discovery
- Administrative Permissions for Scheduling Tasks
- Administrative Permissions for Microsoft Servers

## Administrative Permissions for Grid Members

Grid Member Permissions table below lists the tasks admins can perform and the required permissions for Grid members.

### Grid Member Permissions

Tasks	Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	DNS Views	DNS Zones	Networks	DHCP Ranges
Assign member to DNS zones				RW			RW		
Assign member to networks					RW			RW	
Assign member to DHCP ranges									RW
Configure member properties	RW								
Add a member to a Match Members list of a DNS view	RW								
Delete a view with members in a Match Members list						RO			
View DNS and DHCP member properties		RO	RO						
View and download syslog	RO								
View DNS and DHCP configuration file		RO	RO						



Tasks	Grid Member(s)	Member DNS Properties	Member DHCP Properties	Restart Member DNS	Restart Member DHCP	DNS Views	DNS Zones	Networks	DHCP Ranges
View network statistics	RO								
Restart DNS service on the member				RW					
Restart DHCP service on the member					RW				

### Administrative Permissions for Network Discovery

Limited-access admin groups can initiate a discovery and manage discovered data based on their administrative permissions.

You can set global permissions for network discovery as described in [Defining Global Permissions](#), see [About Administrative Permissions](#). The following table lists the tasks admins can perform and the required permissions for network discovery.

#### Permissions for Network Discovery

Tasks	Network Discovery	DNS Zones	Networks Selected for Discovery	Port Control
Initiate and control a discovery on selected networks	RW		RW	
View discovered data			RO	
Add unmanaged data to existing hosts, and resolve conflicting IP addresses			RW	
Convert unmanaged data to a host, fixed address, reservation, A record, or PTR record		RW	RW	
Configure device interfaces, provision networks on interfaces, de-provision networks	RW			RW

### Administrative Permissions for Scheduling Tasks

You can schedule tasks, such as adding hosts or modifying fixed addresses, for a future date and time. To schedule tasks, you must first enable the scheduling feature at the Grid level, and then define administrative permissions for admin groups and admin roles. For information, see [Scheduled Tasks](#). Only superusers can enable and disable this feature and grant scheduling permissions to admin groups. Limited-access admin groups can schedule tasks only when they have scheduling permissions.

Superusers can do the following:

- Enable and disable task scheduling at the Grid level

- Grant and deny scheduling permissions to admin groups and admin roles
- Schedule tasks for all supported object types
- Reschedule and delete any scheduled task

You can set global permissions to schedule tasks as described in [Defining Global Permissions](#), see [About Administrative Permissions](#). The following table lists the tasks admins can perform and the required permissions. Users with read/write permission to scheduling can view, reschedule, and delete their own scheduled tasks.

*Table 4.10 Scheduling Task Permissions*

Tasks	Scheduling Task	All Networks	All DNS Views	All Shared Record Groups
Schedule the addition, modification, and deletion of all supported object types	RW	RW	RW	RW
View, reschedule, and delete scheduled tasks	RW	RW	RW	RW
Convert unmanaged data to a host, fixed address, reservation, A record, or PTR record	RW	RW	RW	

To schedule tasks for specific resources, admins must have Read/Write permission to scheduling tasks, plus the required permissions to the supported resources. For information about permissions for specific resources, see the following:

- Grid members—See [Administrative Permission for the Grid](#).
- DNS resources—See [Administrative Permissions for DNS Resources](#).
- DHCP resources—See [Administrative Permissions for DHCP Resources](#).

Note that the appliance deletes all pending scheduled tasks when superusers disable task scheduling at the Grid level. The appliance deletes an admin's scheduled tasks when superusers do the following:

- Set the scheduling permission of admin groups and roles to "Deny"
- Delete or disable an admin group or an admin role
- Delete or disable local admins
- Delete the scheduling permission from any admin group or admin role that contains users with pending scheduled tasks
- Change the admin group of a limited-access admin

### Administrative Permissions for Microsoft Servers

By default, only superusers can add Microsoft servers as managed servers to the Grid. Limited-access admins can add and manage Microsoft servers from the Grid based on their administrative permissions. The following table lists the tasks admins can perform and the required permissions. Note that only superusers can add a Microsoft server to a name server group.

#### *Microsoft Server Permissions*

Tasks	Microsoft Server(s)	Grid Member(s)	Network Views	DNS Views	DNS Zones	Resource Records	Networks	DHCP Ranges	Superscopes
Assign Microsoft server to member	RW	RW							
Assign a network view to the Microsoft server	RW	RW	RW						
Assign a DNS view to the Microsoft server	RW	RW		RW					
Assign Microsoft server as primary or secondary to DNS zones	RW			RW	RW				
Remove a Microsoft server as the primary or secondary server of a zone					RW				
Remove a zone from a Microsoft server					RW				
Edit zones and resource records of Microsoft servers					RW	RW			
Assign a Microsoft server to a network	RW						RW		
Assign a Microsoft server to a DHCP range	RW							RW	
Remove a network served by a Microsoft server	RW						RW		
Remove a DHCP range (scope) from a Microsoft server							RW	RW	
Add, modify and remove Microsoft superscopes	RW							RW	RW
Clear leases from Microsoft server	RW							RW	
Edit Microsoft server properties	RW								
View Microsoft server properties	RO								
View and download Microsoft logs	RO								
Start/Stop DNS or DHCP on the Microsoft server	RW								
Remove a Microsoft server from the Grid	RW								

## Administrative Permissions for IPAM Resources

Limited-access admin groups can access certain IPAM resources only if their administrative permissions are defined. By default, the appliance denies access when a limited-access admin group does not have defined permissions. You can grant admin groups read-only or read/write permission, or deny access to the following IPAM resources:

- Network views
- IPv4 networks
- IPv6 networks
- Hosts

The appliance applies permissions for IPAM resources hierarchically. Permissions to a network view apply to all networks and resources in that view. You can also grant an admin group broad permissions to IPAM resources, such as read/write permission to all IPv4 networks and IPv6 networks in the database. In addition, you can grant permission to a specific host in a network. Permissions at more specific levels override global permissions.

The following sections describe the types of permissions that you can set for IPAM resources:

- For more information about [Administrative Permissions for Network Views](#), see [Administrative Permissions for DHCP Resources](#).
- [Administrative Permissions for IPv4 and IPv6 Networks](#), see below
- [Administrative Permissions for Hosts](#), see below.

### Administrative Permissions for IPv4 and IPv6 Networks

Limited-access admin groups can access IPv4 and IPv6 networks only if their administrative permissions are defined. Permissions for a network apply to all its DNS and DHCP resources, if configured. To override network-level permissions, you must define permissions for specific objects within the networks. You can also define permissions for specific DHCP objects and Grid member to restrict admins to perform only the specified DHCP tasks on the specified member. For more information about Defining DNS and DHCP Permissions on Grid Members, see [About Administrative Permissions](#).

You can grant read-only or read/write permission, or deny access to networks, as follows:

- All IPv4 or IPv6 networks—Global permission that applies to all networks in the database.
- A specific network—Network permissions apply to all objects in the network. This overrides global permissions.
- A specific network on a specific member—Network permissions apply to all objects in the network and member permissions apply to the specific member. For information about member permissions and modifying permissions on a Grid Member, see [About Administrative Permissions](#).

### Administrative Permissions for Hosts

A host record can contain both DNS and DHCP attributes if you configure them. When applying administrative permissions to host records, the permissions apply to all relevant DNS and DHCP resources within the host records. You can define global permissions to all hosts. To override global permissions, you must define permissions for specific hosts.

You can grant read-only or read/write permission, or deny access to host records, as follows:

- All hosts—Global permission that applies to all host records in the Grid.
- A specific host—Object permission that applies only to a selected host.

### Administrative Permissions for DHCP Fingerprint Permissions

NIOS provides a global permission for all All DHCP Fingerprints; however, it does not support object level permissions for fingerprints. To use fingerprint filters, you must have superuser privileges.

## Administrative Permissions for Network Insight Tasks

The below table summarizes the permissions you need to perform various tasks related to device discovery.  
*Permissions for Network Discovery*

Tasks	Net work Discovery	DNS Zones	Networks Selected for Discovery	Port Control	All Network Views/ All IPv4 Networks/ All IPv6 Networks	Permissions for Object
Initiate and control a discovery on selected networks	RW		RW			
View discovered data			RO			
Resolve conflicting IP addresses			RW			
Convert unmanaged objects to a host, fixed address, reservation, A record, or PTR record		RW	RW			
Configure device interfaces, provision networks on interfaces	RW			RW		
Configure a Blackout schedule for networks or DHCP ranges	RW				RO	
Creating/editing port reservations for a Grid member, host, fixed address, reservation, A record, or PTR record				RW		RO

## Administrative Permissions for DNS Resources

You can grant roles and admin groups read-only or read/write permission, or deny access to the following DNS resources:

- DNS Views
- DNS Zones
- Response Policy Zones
- All RPZ Rules
- Hosts
- Bulk Hosts
- A records
- AAAA records
- CNAME records
- DNAME records
- MX records
- PTR records
- SRV records
- TXT records

- Hosts
- Bulk Hosts
- Shared Record Groups
- Shared A records
- Shared AAAA records
- Shared CNAME records
- Shared MX records
- Shared SRV records
- Shared TXT records
- DNS64 synthesis groups
- Adding a blank A/AAAA record

The appliance applies permissions for DNS resources hierarchically. Permissions to a DNS view apply to all zones and resource records in that view. Permissions for a zone apply to all its subzones and resource records, and resource record permissions apply to those resource records only. To override permissions set at higher level, you must define permissions at a more specific level. To assign permissions, see [About Administrative Permissions](#).

You can also define permissions for specific DNS objects and Grid member to restrict admins to perform only the specified DNS tasks on the specified member. For more information about defining DNS and DHCP permissions on Grid Members, see [About Administrative Permissions](#).

The following sections describe the different types of permissions that you can set for DNS resources:

- Administrative Permissions for DNS Views
- Administrative Permissions for Zones
- Administrative Permissions for Resource Records

### Administrative Permissions for DNS Views

Limited-access admin groups can access DNS views, including the default view, only if their administrative permissions are defined. Permissions to a DNS view apply to all its zones and resource records. To override view-level permissions, you must define permissions for its zones and resource records. For example, you can grant an admin group read-only permission to a view and read/write permission to all its zones. This allows the admins to display the view properties, but not edit them, and to create, edit and delete zones in the view.

You can grant read-only or read/write permission, or deny access to DNS views, as follows:

- All views—Global permission that applies to all DNS views in the database.
- A specific view—Applies to its properties and its zones, if you do not define zone-level permissions. This overrides the global view permissions.
- All zones in a view—If you do not define permissions for zones, they inherit the permissions of the view they are in.

For information on setting permissions for a view and its zones, see [About Administrative Permissions](#).

The following table lists the tasks admins can perform and the required permissions for DNS views.

#### *Permissions for DNS Views*

Tasks	Grid Member(s)	All DNS Views	Specific DNS View	All DNS Zones
Create, modify, and delete DNS views		RW		
Create, modify, and delete DNS zones with assigned members	RW			RW
Create, modify, and delete DNS zones without assigned members				RW
Modify and delete a specific DNS view			RW	

Tasks	Grid Member(s)	All DNS Views	Specific DNS View	All DNS Zones
Create, modify, and delete DNS zones, subzones, and resource records in a specific DNS view			RW	RW
Add Grid members to a Match Members list of a DNS view	RW		RW	
Delete a DNS view with Grid members in a Match Members list	RW		RW	
View DNS view properties, DNS zones, and resource records		RO		
View DNS zone properties, subzones, and resource records				RO
Restart services from the DNS tab	RO		RW	

### Administrative Permissions for Zones

By default, zones inherit administrative permissions from the DNS view in which they reside. You can override view-level permissions by setting permissions for specific zones. Permissions set for a zone are inherited by its subzones and resource records. To override zone-level permissions, set permissions for specific subzones and resource records. For example, you can grant an admin group the following permissions:

- Read-only to a zone and to all its A, AAAA, and PTR records (in reverse and forward-mapping zones)
- Read/Write permission to all MX and SRV records in the zone
- Deny to all the other resource records—CNAME, DNAME, TXT, host, and bulk host You can grant read-only or read/write permission, or deny access to zones as follows:
- All zones —Global permission that applies to all zones in all views.
- All zones in a view—Permissions at this level override the global permissions.
- A specific zone—Applies to the zone properties and resource records, if you do not define permissions for its resource records. This overrides global and view-level permissions. If you delete a zone and reparent its subzone, the subzone inherits the permissions of the new parent zone.
- All Response Policy Zones—Global permission that applies to all the Response Policy Zones.
- All Response Policy Rules—Global permission that applies to all the local Response Policy Zone rules.

#### Note

Object permissions are not applicable to Response Policy Zone rules.

- Each resource record type in a zone—For example, you can define permissions for all A records and for all PTR records in a zone. If you do not define permissions for resource records, they inherit the permissions of the zone in which they reside.

For information on setting permissions for zones and resource records, see [About Administrative Permissions](#). The following table lists the tasks admins can perform and the required permissions for zones.

#### DNS Zone Permissions

Tasks	Grid Member(s)	Specific DNS View	All DNS Zones	Specific DNS Zone	Resource Records	Shared Record Group
Create, modify, and delete zones, subzones and resource records with assigned members	RW		RW			

Tasks	Grid Member(s)	Specific DNS View	All DNS Zones	Specific DNS Zone	Resource Records	Shared Record Group
Create, modify, and delete zones, subzones and resource records without assigned members			RW			
Lock and unlock a zone				RW		
Delete a zone with assigned Grid members	RW			RW		
Create, modify, and delete all zones, subzones, and resource records in a specific view		RW	RW			
Assign a name server group (member) to a zone	RW			RW		
Delete a zone with name server groups assigned	RW			RW		
Assign a shared record group to a zone				RW		RW
View zone properties, subzones, and resource records of a specific zone				RO		
Search for zones, subzones, and resource records in a specific DNS view		RO	RO			
Copy resource records from one zone to another: Source zone				RO	RO	
Copy resource records from one zone to another: Destination Zone				RW	RW	

## Administrative Permissions for Resource Records

Resource records inherit the permissions of the zone to which they belong. You can override zone-level permissions by setting permissions for specific resource records.

You can grant read-only or read/write permission, or deny access to resource records as follows:

- Each resource record type in all zones and in all views—Global permission that applies to all resource records of the specified type; for example, all A records in the database.
- Each resource record type in a zone—Permissions at this level override global permissions.
- A specific resource record—Overrides zone-level permissions.

For information on setting permissions for resource records, see [About Administrative Permissions](#). The following table lists the tasks admins can perform and the required permissions for resource records.

### *DNS Resources*



Tasks	Resource Record Type	Specific Resource Record
Create, modify, and delete resource records for a specified type, such as all A records or all PTR records	RW	
View resource records for a specified type only	RO	
Search for records of a specified type	RO	
View a specific resource record		RO
View, modify, and delete a specific resource record		RW

The following are additional guidelines:

- Only admins with read/write permission to bulk host records and read/write permission to reverse zones can create bulk host records and automatically add reverse-mapping zones.
- To create host records, admins must have read/write permission to the network and zone of the host.
- Admins must have read-only permission to the host records in a zone to view the Host Name Compliance Report. Admins must have read/write permission to the resource records in a zone to modify host names that do not comply with the host policy.

### Administrative Permissions for Adding Blank A or AAAA Records

By default, only superusers can add and edit A, AAAA, shared A, and shared AAAA records with a blank name. Limited-access admin groups can add and edit A, AAAA, shared A, and shared AAAA records with a blank name, only if their administrative permissions are defined. You can grant read/write or deny permission to **Adding a blank A/AAAA record** for specific admin groups, which applies to all admin roles in the group. You can define global permissions for specific admin groups and roles to allow limited-access users to add and edit blank A, AAAA, shared A, and shared AAAA records, as described in [Defining Global Permissions](#).

### Administrative Permissions for Shared Record Groups

By default, only superusers can add, edit, and delete shared record groups. Limited-access admin groups can access shared record groups, only if their administrative permissions are defined.

You can set different permissions for a shared record group and for each type of shared resource record in the group. For example, you can grant a role or an admin group the following permissions:

- Read-only to a shared record group and to all its shared A, AAAA, and CNAME records
- Read/Write permission to all the shared MX and SRV records in the shared record group
- Deny to the TXT records

You can grant read-only or read/write permission, or deny access to shared record groups, as follows:

- All shared record groups—Global permission that applies to all shared record groups in the database.
- A specific shared record group—Overrides global permissions.
- Each shared record type in all shared record groups — The shared resource record types include shared A records, shared AAAA records, shared CNAME records, shared MX records, shared SRV records, and shared TXT resource records.
- Each shared record type in a shared record group— Permissions at this level override global permissions.
- A specific shared record—Overrides zone-level permissions. Note the following guidelines:
- Shared record group permissions override zone permissions.
- Even if a zone is locked, superusers and limited-access users with read/write access can still edit or delete a shared record in the zone.

For information on setting permissions for shared record groups, see [. The following table lists the tasks admins can perform and the required permissions for shared record groups.](#)

### Permissions for Shared Record Groups

Tasks	All Shared Record Groups	Specific Shared Record Group	Shared Record Type	Specific DNS Zone	Specific Shared Record
Create, modify, and delete shared record groups	RW				
Modify and delete a shared record group		RW			
View a shared record group		RO			
Create, modify, and delete shared records for a specific type			RW		
View or search for shared records of a specific type			RO		
Create, modify, and delete shared records for a specific type in a specified shared record group		RW	RW		
View shared records for a specific type in a specified shared record group only		RO	RO		
Create, modify, and delete a shared record					RW
View a specific shared record					RO
Assign a shared record group to DNS zones		RW		RW	
Change the DNS zones associated with a shared record		RW		RW	
Delete zones with a shared record group assigned. Before you delete a shared record group, you must remove all zones associated with it.		RW		RW	

### Administrative Permissions for DNS64 Synthesis Groups

By default, only superusers can add, edit, and delete DNS64 synthesis groups. Limited-access admin groups can access synthesis groups, only if their administrative permissions are defined.

You can grant read-only or read/write permission, or deny access to synthesis groups, as follows:

- All synthesis groups—Global permission that applies to all shared record groups in the database.
- A specific synthesis group—Overrides global permissions.

For information on setting permissions for synthesis groups, see [About Administrative Permissions](#). The following table lists the tasks admins can perform and the required permissions for synthesis groups.

### Permissions for DNS64 Synthesis Groups

Tasks	All Synthesis Groups Specific Synthesis Group	Specific Shared Record Group	Grid	Specific Member	Specific DNS View
Create, modify, and delete synthesis groups	RW				

Tasks	All Synthesis Groups Specific Synthesis Group	Specific Shared Record Group	Grid	Specific Member	Specific DNS View
Modify and delete a specific synthesis group		RW			
View a synthesis group		RO			
Apply a synthesis group to the Grid		RO	RW		
Apply a synthesis group to a member		RO		RW	
Apply a synthesis group to a DNS view		RO			RW

## Administrative Permissions for DNS Resources with Associated IP addresses in Networks and Ranges

You can further control permissions for DNS resources that have associated IP addresses in a network container, network, or address range. These DNS resources include A records, AAAA records, PTR records, and DNS hosts. Permissions for these resources have been added so you now have more control over who can perform which tasks for these DNS resources without affecting permissions defined for the networks and ranges to which the resources belong. For example, if you want to allow an admin to add, modify, and delete A records associated with IP addresses within a specific network but you do not want the same admin to modify or delete the network, you can grant the admin read-only permission for the specified network and read/write permission for A records in that network. Similar behavior applies to AAAA records, PTR records, and DNS hosts.

As a superuser, you can now grant permissions to admin groups for more granular access to the following resources:

- IPv4 and IPv6 DHCP fixed addresses and IPv4 reservations in a range
- IPv4 and IPv6 host addresses in a range
- A and AAAA records in a network container, network, or range
- IPv4 and IPv6 PTR records in a network container, network, or range

For information about how to configure new permissions for these resources, see [Configuring Permissions for DNS Resources in Networks and Ranges](#) below.

### Best Practices for Configuring Permissions in Networks and Ranges

Before using permissions for DNS resources in networks and ranges, consider the following:

- You can enable and disable these permissions using the `set dns_perm_for_nw_range` CLI command. When you disable permissions after you have enabled and defined them, the appliance retains the permissions in an inactive mode. Inactive permissions are not verified nor displayed in Grid Manager. When you re-enable the permissions, the appliance activates them and displays them in Grid Manager. You can also use the `show dns_perm_for_nw_range` CLI command to verify the status of new permissions.

Note that permissions for fixed addresses and reservations are not controlled by the CLI command; they are always enabled.

You can also enable and disable permissions for DNS Resources in Networks and Ranges through Grid Manager, as described in [Enabling Permissions for DNS Resources in Networks and Ranges](#) below.

- When you switch between enabling and disabling these permissions, changes take effect immediately and a service restart in Grid Manager is not required. However, you may need to refresh Grid Manager to view the changes.

- You can assign these permissions when DNS, DHCP, or Microsoft Management licenses are installed. If you remove all of these licenses after you have configured relevant permissions for supported resources, the appliance retains the permissions, but you will not be able to see the permissions nor configure them.

## Changes to Default Behavior

This section describes changes to the default behavior when you enable permissions for DNS resources with associated IP addresses in networks and ranges. The following table lists behavior in previous releases and changes made in this release for supported resources. Review these changes before you configure permissions for these resources.

Resources	Behavior in Previous NIOS Releases	Changes in NIOS 6.10.4
A Records AAAA Records PTR Records DNS Hosts	<ul style="list-style-type: none"> <li>• Admins can add, modify, and delete A, AAAA, PTR records and DNS hosts that have associated IP addresses in a network or range when they have read/write permission for the respective zone or a higher level DNS parent object (even if they have deny or read-only permission for the network to which the DNS resources belong).</li> </ul>	<ul style="list-style-type: none"> <li>• When you enable new permissions, you can define the following permissions for the admins to add, modify, and delete A, AAAA, PTR records and DNS hosts that have associated IP addresses in a network container, network, or range:               <ul style="list-style-type: none"> <li>• Read/write permission for the specific records in the zone or a higher level DNS parent object.</li> <li><b>and</b></li> <li>• Read/write permission for the records in the specified network container, network, or range to which the resources belong.</li> </ul> </li> </ul> <hr/> <p><b>Note:</b> Fields for A, AAAA, PTR records and DNS hosts in a zone or a higher level DNS parent object, except <b>Name</b>, <b>IP Address</b>, <b>MAC Address</b>, <b>DUID</b> and <b>Disabled</b>, can be modified by admins who do not have write permission for the same records in the specified network container, network, or range.</p> <hr/> <ul style="list-style-type: none"> <li>• You cannot define read-only permission for A, AAAA, and PTR records in a range. The read operation for A, AAAA, PTR, and host records is based on DNS permission hierarchy. In other words, read-only permission for a parent DNS zone allows you to view DNS objects, regardless of the permission you have defined for these resources in their associated ranges.</li> </ul>

Resources	Behavior in Previous NIOS Releases	Changes in NIOS 6.10.4
DHCP-enabled Hosts	<ul style="list-style-type: none"> <li>Admins can add, modify, and delete DHCP-enabled host addresses when they have read/write permission for hosts in the specified zone and read/write permission for the network to which the IP addresses belong. (This behavior stays the same in NIOS 6.10.4.)</li> </ul>	<ul style="list-style-type: none"> <li>When you enable new permissions and you want to allow the admin to add, modify, and delete DHCP-enabled hosts that fall within a specific address range, define read/write permission for hosts in that specified range. Note that if the admin has read/write permission for the network, they can add, modify, and delete hosts that do not fall within a specific address range.</li> </ul> <hr/> <p><b>Note:</b> Fields for DHCP-enabled host addresses, except <b>Name</b>, <b>Enable in DNS</b>, <b>IPv4/IPv6 Address</b>, <b>MAC Address</b>, <b>DUID</b>, and <b>Disabled</b>, can be modified by admins who do not have write permission for the same addresses in the specified network container, network, or range.</p> <hr/> <ul style="list-style-type: none"> <li>Regardless of whether new permissions are enabled, DHCP-enabled host addresses always appear in their respective network containers or networks. However, host addresses in a range appear in the range only when new permissions are enabled.</li> </ul>
Fixed Addresses/ Reservations	<ul style="list-style-type: none"> <li>Admins can add, modify, and delete fixed addresses and reservations in an address range when they have read/write permission for DHCP fixed addresses for the network to which the range belongs.</li> </ul>	<ul style="list-style-type: none"> <li>When you enable new permissions and you want to allow the admin to add, modify, and delete fixed addresses and reservations in a specific address range, define read/write permission for fixed addresses or reservations in that specified range. Note that if the admin has read/write permission for the network, they can add, modify, and delete fixed addresses and reservations that do not fall within a specific address range.</li> <li>Fixed addresses appear in their respective network containers, networks, and ranges regardless of whether new permissions are enabled.</li> </ul>

## Enabling Permissions for DNS Resources in Networks and Ranges

To enable permission for DNS Resources in Networks and Ranges:

- From the **Grid** tab, select the **Grid Manager** tab.
- Expand the Toolbar and select **Grid Properties** -> **Edit**.
- In the *Grid Properties* editor, select the **General** tab -> **Advanced** tab, and then complete the following:
  - Enable DNS Object Permissions in Networks and Ranges:** Select this checkbox to enable DNS object permissions in networks and ranges. When you enable this, admins with Read/Write permission for specific records in a zone or a higher-level DNS parent object, and admins with Read/Write permission for

resource records in specified network containers, networks, or ranges can add, modify, and delete A, AAAA, PTR records, and DNS hosts that have associated IP addresses in the network containers, networks, or ranges.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring Permissions for DNS Resources in Networks and Ranges

To define permissions for resources that are associated with IP addresses in a network container, network, or address range, complete the following:

1. Log in to the appliance through the Infoblox CLI and use the `set dns_perm_for_nw_range` command to enable new permissions, as follows:

```
Infoblox > set dns_perm_for_nw_range on
```

You can also enable permission for DNS resources in networks and ranges through the Infoblox GUI, as described in [Enabling Permissions for DNS Resources in Networks and Ranges](#).

2. Log in to Grid Manager and depending on which permission you want to define, do one of the following:
  - Network View:** From the **Administration** tab, select the **Networks View** tab -> `network_view` checkbox and click the Edit icon.
  - Network Container:** From the **Data Management** tab, select the **IPAM** tab -> `network_container` checkbox and click the Edit icon.
  - Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> `network` checkbox, and then click the Edit icon.
  - DHCP Range:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> **network** -> `addr_range` checkbox, and then click the Edit icon.
  - Fixed Address:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> **network** -> `fixed_address` checkbox, and then click the Edit icon.
  - Reservation:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> **network** -> `reservation` checkbox, and then click the Edit icon.
  - Zone:** From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> `zone` checkbox, and then click the Edit icon.

Note that you cannot assign permissions for zones that are auto-created.
3. In the editor, click the **Permissions** tab, and select the supported permission from the **Permissions** drop-down list for the admin group or role.
4. Select a resource from the drop-down list in the **Resources** column.
5. Save the configuration.

## Permission Examples

The following table lists examples for configuring new permissions for fixed addresses (or reservations) in network 10.1.2.0/24 and range 10.1.2.1-10.1.2.10.

Action	Permission for network 10.1.2.0/24	Permission for range 10.1.2.1-10.1.2.10	Action Allowed? (Yes/No)	Comment
Add, modify, or delete fixed address 10.1.2.5	No	No	No	N/A
Add, modify, or delete fixed address 10.1.2.5	No	Read/write for "Fixed addresses in 10.1.2.1-10.1.2.10 Range"	Yes	Read/write permission at the range level is sufficient for creating a fixed address that falls within the range.

Action	Permission for network 10.1.2.0/24	Permission for range 10.1.2.1-10.1.2.10	Action Allowed? (Yes/No)	Comment
Add, modify, or delete fixed address 10.1.2.100	Read/write for "Fixed addresses in 10.1.2.0/24 Network"	Deny for "Fixed addresses in 10.1.2.1-10.1.2.10 Range"	Yes	Since fixed address 10.1.2.100 does not belong to the 10.1.2.1-10.1.2.10 range, read/write permission for "Fixed addresses in 10.1.2.0/24 Network" is sufficient for the operation.

The following table lists some examples for configuring DNS resources that have associated IP addresses in a network or range:

Action	Permission for DNS zone corpxyz.com	Permission for network 10.1.2.0/24	Permission for range 10.1.2.1-10.1.2.10	Action Allowed? (Yes/No)	Comment
Add, modify, or delete an A record with IP address 10.1.2.8	Read/write permission for corpxyz.com	No	No	No	Read/write permission for "A Records in 10.1.2.1-10.1.2.10 range" is also required.
Add, modify, or delete an A record with IP address 10.1.2.8	Read/write permission for corpxyz.com	No	Read/write for "A Records in 10.1.2.1-10.1.2.10	Yes	Since 10.1.2.8 falls within the 10.1.2.1-10.1.2.10 range, read/write permission for "A Records in 10.1.2.1-10.1.2.10 Range" and read/write for corpxyz.com are both required.
Add, modify, or delete an A record with IP address 10.1.2.8, and modify or delete a network	Read/write permission for corpxyz.com	Read-only permission for network 10.1.2.0/24	Read/write for "A Records in 10.1.2.1-10.1.2.10 Range	Yes for A record No for network	Admins can add, modify, and delete A records because they have read/write permissions for the zone and range; but they cannot modify or delete networks because they have read-only permission for network 10.1.2.0/24.
Add, modify or delete DHCP-enabled host address 10.1.2.22	Yes if the host is a DNS host. N/A if the host is a DHCP host.	Read/write permission for "IPv4 Hosts in 10.1.2.0 network"	No	Yes	Host address 10.1.2.22 is within the 10.1.2.0 network but outside of the 10.1.2.1-10.1.2.10 range, so read/write permission for "IPv4 Hosts in 10.1.2.0 network" is sufficient.
Add, modify, or delete DHCP-enabled host address 10.1.2.8, and modify or delete a network	Yes if the host is a DNS host. N/A if the host is a DHCP host.	Read-only permission for network 10.1.2.0/24	Read/write for "Hosts in 10.1.2.1-10.1.2.10 Range	Yes for A record No for network	Admins can add, modify, and delete DHCP-enabled hosts because they have read/write permissions for "Hosts in 10.1.2.1-10.1.2.10 range"; but they cannot modify or delete networks because they have read-only permission for network 10.1.2.0/24.

The following table list an example for permissions required to configure PTR records that have associated IP addresses in a network:

Action	Permission for DNS zone corpxyz.com	Permission for network 10.1.2.0/24	Permission for reverse zone 0.0.10.in-addr.arpa	Action Allowed? (Yes/No)	Comment
Add, modify, or delete a PTR record with IP address 5.0.0.10. <b>Note:</b> You can also add, modify, or delete PTR records in the IPv6 reverse-mapping zone.	Read/write permission for corpxyz.com	No	Yes	Yes	Read/write permission for "PTR Records in corpxyz.com and 0.0.10.in-addr.arpa" is required.

## Administrative Permissions for DHCP Resources

Limited-access admin groups can access certain DHCP resources only if their administrative permissions are defined. By default, the appliance denies access when a limited-access admin group does not have defined permissions. You can grant admin groups read-only or read/write permission, or deny access to the following DHCP resources:

- Network views
- IPv4 networks
- Hosts
- IPv4 DHCP ranges
- IPv4 DHCP fixed addresses
- IPv4 DHCP reservations
- MAC address filters
- IPv4 shared networks
- IPv4 network templates
- IPv4 DHCP range templates
- IPv4 fixed address templates
- IPv4 DHCP enabled host addresses
- IPv4 DHCP lease history
- Roaming hosts
- IPv6 networks
- IPv6 DHCP ranges
- IPv6 DHCP fixed addresses
- IPv6 DHCP enabled host addresses
- IPv6 shared networks
- IPv6 network templates
- IPv6 DHCP range templates
- IPv6 fixed address templates
- IPv6 DHCP lease history

You can grant an admin group broad permissions to DHCP resources, such as read/write permission to all IPv4 or IPv6 networks and shared networks in the database. In addition, you can grant permission to specific resources, such as a specific IPv4 or IPv6 network or DHCP range, or an individual address in an IPv4 or IPv6 network. Permissions at more specific levels override global permissions.

You can also define permissions for specific DHCP objects and Grid member to restrict admins to perform only the specified DHCP tasks on the specified member. For more information about defining DNS and DHCP permissions on Grid Members, see [About Administrative Permissions](#).

The following sections describe the different types of permissions that you can set for DHCP resources:

- [Administrative Permissions for Network Views](#)
- [Administrative Permissions for IPv4 and IPv6 Networks and Shared Networks](#)
- [Administrative Permissions for IPv4 or IPv6 Fixed Addresses and IPv4 Reservations](#)
- [Administrative Permissions for IPv4 and IPv6 DHCP Ranges](#)
- [Administrative Permissions for IPv4 or IPv6 DHCP Templates](#)
- [Administrative Permissions for Roaming Hosts](#)
- [Administrative Permissions for MAC Address Filters](#)



- Administrative Permissions for the IPv4 and IPv6 DHCP Lease Histories

## Administrative Permissions for Network Views

Limited-access admin groups can access network views, including the default network view, only if they have read-only or read/write permission to a specific network view or to all network views. Permissions granted to a network view apply to all its IPv4 and IPv6 networks, shared networks, DHCP ranges and fixed addresses.

You can grant admin groups read-only or read/write permission, or deny access to network views as follows:

- All network views—Global permission that applies to all network views in the database.
- A specific network view—Permission to a specific network view applies to the properties you set in the *Network View* editor, and to all the IPv4 and IPv6 networks and shared networks in the network view. This overrides the global permission to all network views. When you configure permissions for a network view, you can also set permissions for the following:
  - All IPv4 and IPv6 networks in the selected network view—If you do not define permissions for IPv4 or IPv6 networks, they inherit the permissions of their network view.
  - All IPv4 and IPv6 shared networks in a specific network view—If you do not define permissions for IPv4 or IPv6 shared networks, they inherit the permissions of their network view.

Note that you can grant an admin group read-only or read/write permission to specific IPv4 or IPv6 networks in a network view, without granting them permission to that network view. For information, see [Permissions for IPv4 and IPv6 Networks and Shared Networks](#) below.

For information on how to define permissions for network views, see [About Administrative Permissions](#).

The following table lists the tasks admins can perform and the required permissions for network views.

### Network View Permissions

Tasks	All DNS Views	Specific DNS View	All Network Views	Specific Network View	All IPv4 or IPv6 Networks	All IPv4 or IPv6 Shared Networks
Create and delete network views and their associated DNS views	RW		RW			
Create and delete a network view and its associated DNS views		RW		RW		
Create, modify, and delete IPv4 and IPv6 networks and shared networks in all network views			RW			
Create, modify, and delete IPv4 and IPv6 networks and shared networks in a network view				RW		

Tasks	All DNS Views	Specific DNS View	All Network Views	Specific Network View	All IPv4 or IPv6 Networks	All IPv4 or IPv6 Shared Networks
View the properties of all network views			RO			
View network statistics of all network views			RO			
View and search for all IPv4 and IPv6 networks and shared networks			RO			
View the properties of a network view				RO		
View and search for IPv4 and IPv6 networks and shared networks in a network view				RO		
Expand and join IPv4 and IPv6 networks			RW			
Expand and join IPv4 and IPv6 networks in a specific network view				RW		
Create, modify, and delete IPv4 and IPv6 networks, DHCP ranges and fixed addresses in a specific network view				RW		
View network statistics and properties of all networks in a network view				RO		
Search for IPv4 and IPv6 networks in a network view				RO		

Tasks	All DNS Views	Specific DNS View	All Network Views	Specific Network View	All IPv4 or IPv6 Networks	All IPv4 or IPv6 Shared Networks
Create, modify, and delete all IPv4 or IPv6 shared networks						RW
View the properties of all IPv4 or IPv6 shared networks						RO
View and search for IPv4 and IPv6 shared networks in a network view				RO		
Restart services from the DHCP tab	RO			RW		

### Administrative Permissions for IPv4 and IPv6 Networks and Shared Networks

Limited-access admin groups can access IPv4 and IPv6 networks, including shared networks, only if their administrative permissions are defined. Permissions for a network apply to all its DHCP ranges and fixed addresses. To override network-level permissions, you must define permissions for specific DHCP ranges and fixed addresses. For example, you can grant an admin group read-only permission to a network, read/write permission to its DHCP ranges, and read-only permission to its fixed addresses.

You can grant read-only or read/write permission, or deny access to networks, as follows:

- All IPv4 or IPv6 networks—Global permission that applies to all IPv4 or all IPv6 networks in the database.
- All IPv4 or IPv6 shared networks—Global permission that applies to all IPv4 or all IPv6 shared networks in the database.
- A specific IPv4 or IPv6 network—Network permissions apply to its properties and to all DHCP ranges, fixed addresses and hosts in the network, if they do not have permissions defined. This overrides global permissions.
- All IPv4 or IPv6 DHCP ranges in a network—If you do not define permissions for DHCP ranges, they inherit the permissions of the network in which they reside.
- All IPv4 or IPv6 fixed addresses in a network—If you do not define permissions for fixed addresses, they inherit the permissions of the network in which they reside.

To define permissions for a specific IPv4 or IPv6 network and its DHCP ranges and fixed addresses, see [About Administrative Permissions](#).

The following table lists the tasks admins can perform and the required permissions for IPv4 and IPv6 networks.

#### Network Permissions

Tasks	Grid Member(s)	All IPv4 or IPv6 Networks	Specific IPv4 or IPv6 Network	All IPv4 or IPv6 Shared Networks	Specific DNS Zone	All IPv4 or IPv6 DHCP Ranges	All IPv4 or IPv6 Fixed Addresses	IPv4 or IPv6 Network Template
-------	----------------	---------------------------	-------------------------------	----------------------------------	-------------------	------------------------------	----------------------------------	-------------------------------

Create, modify, and delete IPv4 or IPv6 networks, DHCP ranges, and fixed addresses without assigned Grid members		<b>RW</b>						
Create, modify, and delete IPv4 or IPv6 networks, DHCP ranges, and fixed addresses with assigned Grid members	<b>RW</b>	<b>RW</b>						
Assign a Grid member to a specific IPv4 or IPv6 network and its DHCP ranges	<b>RW</b>		<b>RW</b>					
Expand and join IPv4 or IPv6 networks		<b>RW</b>						
Create IPv4 or IPv6 networks from templates		<b>RW</b>						<b>RO</b>
Create, modify, and delete an IPv4 or IPv6 network		<b>RW</b>						
View IPv4 or IPv6 network properties and statistics, and search for DHCP ranges and fixed addresses in a specific network			<b>RO</b>					
Create, modify, and delete IPv4 or IPv6 DHCP ranges and fixed addresses in a specific network			<b>RW</b>					

Create and split an IPv4 or IPv6 network and automatically create a reverse DNS zone			RW		RW			
Create, modify, and delete IPv4 or IPv6 shared networks					RW			
View IPv4 or IPv6 shared networks					RO			
Create, modify, and delete IPv4 or IPv6 DHCP ranges with an assigned member in a specific network	RW		RW					
Create, modify, and delete IPv4 or IPv6 DHCP ranges						RW		
View and search for IPv4 or IPv6 DHCP ranges in a specific network			RO					
Create, modify, and delete IPv4 or IPv6 fixed addresses							RW	
View and search for IPv4 or IPv6 fixed addresses in a specific network			RO					

### Administrative Permissions for IPv4 or IPv6 Fixed Addresses and IPv4 Reservations

IPv4 and IPv6 fixed addresses and IPv4 reservations inherit the permissions of the networks in which they reside. You can override network-level permissions by defining permissions for fixed addresses.

You can grant read-only or read-write permission, or deny access to fixed addresses, as follows:

- All IPv4 fixed addresses/reservations—Global permission that applies to all IPv4 fixed addresses and reservations in the database.
- All IPv6 fixed addresses—Global permission that applies to all IPv6 fixed addresses in the database.
- All IPv4 fixed addresses/reservations in a network—Permissions at this level override global permissions. If you do not define permissions for fixed addresses and reservations, they inherit the permissions of the network in which they reside.

- All IPv6 fixed addresses in a network— Permissions at this level override global permissions. If you do not define permissions for IPv6 fixed addresses, they inherit the permissions of the network in which they reside.
- A single IPv4 fixed address/reservation—Overrides global and network-level permissions.
- A single IPv6 fixed address—Overrides global and network-level permissions.

For information on setting permissions for fixed addresses, see [About Administrative Permissions](#).

The following table lists the tasks admins can perform and the required permissions for IPv4 and IPv6 fixed addresses.

*Permissions for Fixed Addresses/Reservations*

Tasks	Specific IPv4 or IPv6 Network	All IPv4 or IPv6 fixed Addresses/ IPv4 Reservations	Specific IPv4 or IPv6 Fixed Address/ IPv4 Reservation
Create, modify, and delete IPv4 fixed addresses/reservations or IPv6 fixed addresses		<b>RW</b>	
Create, modify, and delete IPv4 fixed addresses/reservations or IPv6 fixed addresses in a specific network	<b>RW</b>		
Modify and delete an IPv4 fixed address/reservation or IPv6 fixed address			<b>RW</b>
View and search for all IPv4 fixed addresses/reservations or IPv6 fixed addresses		<b>RO</b>	
View and search for IPv4 fixed addresses/reservations or IPv6 fixed addresses in a network	<b>RO</b>	<b>RO</b>	
View and search for an IPv4 fixed address/reservation or IPv6 fixed address			<b>RO</b>

**Administrative Permissions for IPv4 or IPv6 DHCP Enabled Host Addresses**

A read-write permission to IPv4 or IPv6 Host Address gives limited-access users the ability to create, modify, and delete IPv4 and IPv6 DHCP enabled host addresses in a specified network. Admin users with a read-write permission can create, modify, and delete IPv4 or IPv6 DHCP enabled host addresses only in the specified network. They do not have the ability to create, modify or delete any networks or objects, such as fixed addresses, in those networks.

You can also grant admin users read-only permission or deny access to the following:

- IPv4 Host Address—Object permission that applies to all IPv4 DHCP enabled host addresses in a specified network.
- IPv6 Host Address—Object permission that applies to all IPv6 DHCP enabled host addresses in a specified network.

For information about setting permissions for DHCP enabled host addresses, see [About Administrative Permissions](#).

The following table lists tasks that admins can perform and the required permissions for IPv4 and IPv6 DHCP enabled host addresses.

*Permissions for DHCP Enabled Host Addresses*

Tasks	Specific IPv4 or IPv6 Network	All IPv4 or IPv6 DHCP enabled host Addresses
Create, modify, and delete IPv4 or IPv6 DHCP enabled host addresses in a specified network		RW
Modify and delete a specific IPv4 or IPv6 DHCP enabled host address		RW
View and search for all IPv4 or IPv6 DHCP enabled host addresses		RO
View and search for IPv4 or IPv6 DHCP enabled host addresses in a specified network		RO

### Administrative Permissions for IPv4 and IPv6 DHCP Ranges

DHCP ranges inherit the permissions of the networks in which they reside. You can override network-level permissions by defining permissions for DHCP ranges. You can read-only or read/write permission, or deny access to DHCP address ranges, as follows:

- All IPv4 or IPv6 DHCP ranges—Global permission that applies to all IPv4 or IPv6 DHCP ranges in the database.
- All IPv4 or IPv6 DHCP ranges in a network—Permissions at this level override global permissions. If you do not define permissions for DHCP ranges, they inherit the permissions of the network in which they reside.
- A single IPv4 or IPv6 DHCP range—Overrides global and network-level permissions.

For information on setting permissions for DHCP ranges, see [About Administrative Permissions](#). The following table lists the tasks admin can perform and the required permissions for DHCP ranges.

#### DHCP Ranges

Tasks	Grid Member(s)	Specific IPv4 or IPv6 Network	All DHCP IPv4 or IPv6 Ranges	Specific IPv4 or IPv6 DHCP Range	MAC Address Filter
Create, modify, and delete IPv4 or IPv6 DHCP ranges with an assigned member or a failover association	RW		RW		
Create, modify, and delete IPv4 or IPv6 DHCP ranges in a network with assigned members	RW	RW			
Modify and delete an IPv4 or IPv6 DHCP range with an assigned member	RW			RW	
View and search for all IPv4 or IPv6 DHCP ranges with an assigned member	RO			RO	
View and search for IPv4 or IPv6 DHCP ranges in a network with assigned members	RO	RO			

Tasks	Grid Member(s)	Specific IPv4 or IPv6 Network	All DHCP IPv4 or IPv6 Ranges	Specific IPv4 or IPv6 DHCP Range	MAC Address Filter
View and search for an IPv4 or IPv6 DHCP range with an assigned member	RO			RO	
View and search for an IPv4 or IPv6 DHCP range without an assigned member				RO	
Apply relay agent and option filters to an IPv4 DHCP range				RW	
Apply a MAC address filter to an IPv4 DHCP range				RW	RO

### Administrative Permissions for IPv4 or IPv6 DHCP Templates

There are three types of DHCP templates for IPv4 and IPv6 objects—network, DHCP range, and fixed address/reservation templates. To access any of these templates, a limited-access admin group must have read-only permission to the template. Limited-access admin groups cannot have read/write permission to the templates. Only superusers can create, modify and delete network, DHCP range, and fixed address templates. An admin group with read-only permission to the DHCP templates can view them and use them to create networks, DHCP ranges and fixed addresses, as long as they have read/write permissions to those DHCP resources as well.

You can set global read-only permission that applies to all DHCP templates, and you can set permissions to specific templates as well.

For information on setting permissions, see [About Administrative Permissions](#). The following table lists the tasks admins can perform and the required permissions for DHCP templates.

#### Permissions for DHCP Templates

Tasks	IPv4 or IPv6 DHCP Templates	All IPv4 or IPv6 Networks	All IPv4 or IPv6 DHCP Ranges	All IPv4 or IPv6 Fixed Addresses/ IPv4 Reservations
Create IPv4 or IPv6 networks from templates	RO	RW		
Create IPv4 or IPv6 DHCP ranges from templates	RO		RW	
Create IPv4 fixed addresses/reservations or IPv6 fixed addresses from templates	RO			RW
View templates	RO			

Note the following additional guidelines:

- DHCP range templates and fixed address templates do not inherit their permissions from network templates. You must set permissions for each type of template.
- An admin group can create a network using a network template that includes a DHCP range template and a fixed address template, even if it has no permission to access the DHCP range and fixed address templates.



## Administrative Permissions for Roaming Hosts

Limited-access admin groups can access roaming hosts only if their administrative permissions are defined. The appliance denies access to roaming hosts for which an admin group does not have defined permissions. You can grant read-only or read/write permission, or deny access to roaming hosts as follows:

- All roaming hosts in the database—Global permission that applies to all the roaming hosts in the database.
- A specific roaming host—Permissions that applies to specific roaming host.

For information on setting permissions, see [About Administrative Permissions](#). The following table lists the tasks admins can perform and the required permissions for roaming host.

### Permissions for Roaming Hosts

Tasks	Grid DHCP Properties	Specific IPv4 or IPv6 Roaming Host	All Roaming Host
Enable roaming hosts	<b>RW</b>		
View roaming host	<b>RO</b>	<b>RO</b>	<b>RO</b>
Create, modify, and delete roaming hosts	<b>RO</b>		<b>RW</b>
Modify and delete roaming host	<b>RO</b>	<b>RW</b>	

## Administrative Permissions for MAC Address Filters

Limited-access admin groups can access MAC address filters only if their administrative permissions are defined. The appliance denies access to MAC address filters for which an admin group does not have defined permissions. You can grant read-only or read/write permission, or deny access to MAC address filters as follows:

- All MAC address filters in the database
- A specific MAC address filter

For information on setting permissions, see [About Administrative Permissions](#). The following table lists the tasks admins can perform and the required permissions for MAC address filters.

### Permissions for MAC Filters

Tasks	All MAC Address Filters	Specific MAC Address Filter	Specific IPv4 DHCP Ranges
Create, modify, and delete MAC address filters	<b>RW</b>		
Create, modify, and delete MAC address entries for a MAC address filter		<b>RW</b>	
Modify and delete a MAC address filter		<b>RW</b>	
Apply a MAC address filter to an IPv4 DHCP range		<b>RO</b>	<b>RW</b>
Delete a MAC address filter from an IPv4 DHCP range		<b>RO</b>	<b>RW</b>

Tasks	All MAC Address Filters	Specific MAC Address Filter	Specific IPv4 DHCP Ranges
View MAC address filters and their MAC address entries	RO		
View a MAC address filter and its MAC address entries		RO	

### Administrative Permissions for the IPv4 and IPv6 DHCP Lease Histories

A limited-access admin group can view and export the IPv4 and IPv6 DHCP lease histories if it has read-only permission to the IPv4 and IPv6 DHCP lease history. Permissions to the IPv4 and IPv6 DHCP lease histories are different from the network permissions. Therefore, an admin group can access the IPv4 and IPv6 DHCP lease histories, regardless of its network permissions. Note that only superusers can import a DHCP lease history file.

To define permissions for the IPv4 and IPv6 DHCP lease histories:

- For an admin group: From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin\_group* in the Groups table, and then click the Add icon -> **Global Permissions** from the Create New Permission area or select Add -> **Global Permissions** from the Toolbar.  
or  
For an admin role: From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin\_role* in the Roles table, and then click Add icon -> **Global Permissions** from the Create New Permission area or select Add -> **Global Permissions** from the Toolbar.
- Complete the following in the *Manage Global Permissions* dialog box:
  - Permission Type:** Select **DHCP Permissions** from the drop-down list.
  - In the table, select **Read/Write**, **Read-only**, or **Deny** for **All IPv4 DHCP Lease History** and **All IPv6 DHCP Lease History**.
- Save the configuration and click **Restart** if it appears at the top of the screen.

### Administrative Permissions for File Distribution Services

You can restrict access to the TFTP, HTTP and FTP services provided by the appliance. By default, the appliance denies access to the TFTP, HTTP and FTP services, unless an admin group has their administrative permissions defined.

You can grant read-only or read/write permission, or deny access to the following resources:

- Grid File Distribution Properties—Applies to the Grid and its members, directories, and files. You can set this from the Administrators perspective only.
- Member File Distribution Properties—Applies to the Grid member properties only.
- A specific directory—Applies to the directory and its files.

For information on setting permissions, see [About Administrative Permissions](#). The following table lists the tasks admins can perform and the required permissions for file distribution services.

#### Permissions for File Distribution Services

Tasks	Grid File Distribution Properties	Member Distribution Properties	Specific Directory
Create and remove directories and files	RW		
Modify the Grid and member file distribution properties	RW		
View the Grid and member file distribution properties, directories, and files	RO		

Tasks	Grid File Distribution Properties	Member Distribution Properties	Specific Directory
Modify the member file distribution properties		RW	
View the member file distribution properties		RO	
Add and delete a directory, subdirectories, and files in the directory			RW
View a directory and its subdirectories and files			RO

## Administrative Permissions for Dashboard Tasks

Limited-access admin groups can configure IPAM tasks on the Tasks Dashboard only if their administrative permissions are defined. The appliance denies access to IPAM tasks for which an admin group does not have defined permissions.

You can grant read-only or read/write permission, or deny access to IPAM tasks as follows:

- All IPAM tasks on the Tasks Dashboard
- A specific IPAM task

When you deny access to an IPAM task for an admin group, users cannot configure the task on their dashboards. Users must have at least read-only permission to a specific task to see it in the task pack. To perform a specific task, users must also have read/write permission to the objects associated with the task. For information about specific permissions for IPAM, DNS, and DHCP objects, see [Administrative Permissions for IPAM Resources](#), [Administrative Permissions for DNS Resources](#), and [Administrative Permissions for DHCP Resources](#). For information about setting permissions, see [About Administrative Permissions](#). The following table lists the tasks admins can perform and the required permissions for configuring IPAM tasks on the Tasks Dashboard.

### Permissions for IPAM Tasks

Tasks	All Dashboard Tasks	Add Networks	Add Hosts	Add Fixed Addresses	Add CNAME Record	Add TXT Record	Add MX Record
Configure all tasks in the IPAM task pack	RO RW						
Configure the Add Networks task		RO RW					
Configure the Add Hosts task			RO RW				
Configure the Add Fixed Addresses task				RO RW			
Configure the Add CNAME Record task					RO RW		

Tasks	All Dashboard Tasks	Add Networks	Add Hosts	Add Fixed Addresses	Add CNAME Record	Add TXT Record	Add MX Record
Configure the Add TXT Record task						RO RW	
Configure the Add MX Record task							RO RW

## Administrative Permissions for Certificate Authentication Services and CA Certificates

Limited-access admins can configure certificate authentication services and CA certificates only if their administrative roles and permissions are defined. If you want to allow admins to configure two-factor authentication, you can assign the PKI Admin role to limited-access admins or grant them read/write permissions to the following:

- All certificate authentication services
- All CA Certificates

For information about setting permissions, see [About Administrative Permissions](#). The following table lists the admin tasks and required permissions for configuring certificate authentication services and managing CA certificates.

### Administration Permissions

Tasks	Grid Member(s)	All Certificate Authentication Services	All CA Certificates
Create, modify, and delete certificate authentication services		RW	
Create, modify, and delete CA certificates	RW		RW

## Administrative Permissions for Object Change Tracking

Infoblox stores updated objects in the NIOS database. Users with read-only permission can query and view these objects. Grid Manager allows you to grant the following permissions:

- Read-only permission to view deleted objects information. If the users have a read-only permission and the **exclude\_deleted** flag is not set, then they will receive information about the deleted objects only.
- Deny to prevent the users from accessing updated objects information. If the users have a Deny permission, then they cannot query for any object updates even if the **exclude\_deleted** flag is not set.

Users with a read-only permission must have permissions on all the objects to perform a full or an incremental synchronization. For example, consider that a user, user1, has read-only permission on DNS views, but does not have permission on network views. If user1 performs a full synchronization, NIOS does not include network views in the response as user1 does not have permission to view these objects. Hence, Infoblox recommends that you give permissions to all the objects in the NIOS database.

## Administrative Permissions for Load Balancers

Limited-access admins can view and manage GLBs (Global Load Balancers), load balancer synchronization groups, and their associated objects if their administrative roles and permissions are defined. If you want to allow admins to manage GLB objects, assign the Load Balancer Admin role to limited-access admins and grant them the following permissions:

- Read/write, read-only, or deny permission to NIOS managed GLB groups and independent load balancers

- Read/write, read-only, or deny permission to NIOS managed GLB objects
- Read-only or deny permission to GLB objects, such as DNS profiles and iRules, that are synchronized from the GLB but cannot be managed through NIOS

For information about setting permissions, see [About Administrative Permissions](#). The following table lists the admin tasks and required permissions for configuring GLBs, load balancer synchronization groups, and their associated objects.

*Administration Permissions for Load Balancers and Load Balancer Synchronization Groups*

Tasks	Grid Member(s)	All Load Balancer Objects	All Load Balancers	All Load Balancer Groups
View load balancer objects	RO	RO		
Add and modify synchronized load balancer objects	RW	RW		
Add, modify, and delete synchronized load balancer objects	RW	RW	RW	RW
Modify and delete synchronized load balancer groups	RW		RW	RW

## Authenticating Admins Using SAML

NIOS uses SAML (Security Assertion Markup Language) 2.0 authentication support for Single-Sign-On in NIOS. SAML provides a standard vendor-independent grammar and protocol for transferring information about a user from one web server to another independent of the server DNS domains. SAML enables IT administrators to manage user access rights in a single place. By enabling SAML, user management is delegated to an external application, thus relieving IT administrators the complexity of maintaining user accounts in all the applications (also known as Service Providers) being used by the organization. Instead, IT administrators need to maintain one account in the Identity Provider (IdP) which can be used across Service Providers (SPs). IdP is the application server that maintains the user accounts of the entire organization. IT administrators can manage users access rights at one place. User can login to the IdP directly and once logged in, they can be traverse towards the required SP without being prompted for the user ID and password. SAML helps NIOS delegate Identity Management to a third-party SSO application (IdP) and thereby eases administrative efforts.



Note

You need super user permissions to perform SAML-related configurations.

## SAML Login Use Cases

The following is a list of use cases and the outcome of NIOS users trying to log in using SAML authentication and not using SAML authentication:

- If SAML is enabled and users have already logged in to the IdP account and the corresponding user account is present in NIOS, users can directly start using Grid Manager without logging in to NIOS.
- If a user has logged in to the IdP account and the corresponding IdP account is not present in NIOS, if the **Auto Create User** checkbox is selected, the user can directly start using Grid Manager without logging in to NIOS. For information about the **Auto Create User** checkbox, see [Auto Creating SAML Users in NIOS](#).
- If a NIOS user who is not SAML-authorized tries to log in to NIOS using the **SSO Login** button, the login fails. However, the user can log in using the **Login** button.

## Prerequisites for Configuring SAML Authentication

Ensure that you meet the following prerequisites before you configure SAML for NIOS:

- When adding the NIOS application in IdP, specify the Grid Manager URL in the `https://<Grid Manager IP address>:8765/?acs` format. This is referred to as the Assertion Consumer Service URL or ACS URL. The 8765 port is opened for SAML services.
- After you add NIOS to the IdP, either copy the metadata or the metadata URL or specify it in the SAML configuration screen.
- Ports 443 (HTTPS) and 80 (HTTP) must be allowed on the firewall to allow NIOS to communicate with IdP.
- Ensure that the group that you specify in the IdP also exists in NIOS with the same users as that in the IdP. If you did not specify a group attribute in IdP, SAML authenticated users are added to the default SAML group: **saml-group**.
- SAML authentication in NIOS requires configuring an Identity Provider (IdP) for authentication. Infoblox-verified named IdPs are listed in the **IDP Type** drop-down list. The **IDP Type** drop-down list also contains the **Others** option for users who wish to configure an IdP that is not listed. Due to the lack of compliance to SAML standards and widely varying IdP vendor implementations, Infoblox is unable to provide configuration support if you select the **Others** option. Infoblox recommends that you contact the IdP vendor for support if you use this option.

## SAML Qualified IDPs and Response Format

The qualified SAML IDPs are:

- Azure SSO
- Ping Identity
- Shibboleth SSO
- OKTA

IdP's response is as following:

For example:

```
<saml2:NameID Format="urn:oasis:names:tc:SAML:1.1:nameid-format:unspecified">arunchurimanjunath@gmail.com</saml2:NameID>
```

NameID is the mandatory parameter in the IDP response format.

## Enabling SAML Authentication

To enable SAML authentication for NIOS users, perform the following steps:

1. Login as a super user.
2. Click the **SAML Authentication Services** tab.
3. From the Toolbar, click **Add** -> **SAML Service**.
4. In the *Add SAML Authentication Service* wizard:
  - **Name**: Enter a name for the authentication service.
  - **IDP Type**: Select the IdP type that you want to configure for the authentication service. NIOS supports the following IdP types: Azure SSO, Okta, Ping Identity, Shibboleth SSO, Others.
  - **SSO Metadata URL**: Enter the metadata URL of the IdP. Alternatively, copy the metadata into a file and upload the file. For information on obtaining the metadata, see [Obtaining Metadata](#).
  - **SSO Redirect IP Address/FQDN**: Enter the IP address or the FQDN of Grid Master. If you enter a value other than the IP address or FQDN, the SAML service will not work. A best practice is to enter the FQDN because it is used by the IdP for the SAML response.
  - **Session Timeout**: Enter the amount of time that a SAML user can be idle after which the session must terminate. The time that you specify here supersedes the session timeout time specified in the **Grid Properties** - > **Security** - > **Session Timeout(s)** field. For more information about session timeout in the Grid, see [Managing Security Operations](#).
  - **IDP Group Attribute**: Enter a group attribute name. The group attribute name that you enter must have the same value as the **Group Attribute** in your IdP configuration. If the value in NIOS is different from what is configured in the IdP, or if the value is missing, the incoming SAML user is placed in the default SAML

group **saml-group**. If the value matches, the IdP's group attribute filter passes the user's group membership to NIOS. If there is a NIOS group that exactly matches a group name from the list of groups from the IdP, and the NIOS group is configured for SAML, the user get placed into that NIOS group. If there is no matching NIOS group, the user gets placed in the default SAML group **saml-group**.

- **Comment:** Enter additional comments if any.

5. Click **Save & Close**.

Now, if you log out and try to relogin, the **SSO Login** button is displayed.

When SAML authorized users try to login for the first time by clicking the **SSO Login** button, they are directed to their IdP login page. If the user credentials they specified exist in the IdP, they are redirected to the NIOS home page.



#### Note

- If there are any IdP or metadata changes following a Grid Master Candidate promotion, it is necessary to delete the configuration and re-configure.
- After HA failover, ensure that you reload the Grid Manager after 3 to 5 seconds, to view the SSO login button.

## Obtaining Metadata

This section explains how to obtain the metadata URL of the IdP. The procedures in this section may vary a little depending on the type of IdP that you select. The procedure in this section uses Okta as an IdP example. If you are using an IdP other than Okta, contact your IT administrator for the metadata URL.

To obtain the metadata URL of Okta:

1. Log in to your Okta account.
2. Go to **My Applications** and click the URL of your Grid Manager.
3. You can either copy the XML metadata for the Grid Manager into a file or use the URL of the metadata.

## Auto Creating SAML Users in NIOS

After enabling SAML authentication, you can configure NIOS such that users who belong to a particular group in IdP are automatically created in NIOS. Once the users are automatically created in NIOS, if they log in to their IdP account, they can directly access Grid Manager.

1. Login as a super user.
2. Create a group by the same name as that of the group in the IdP account. For information about creating groups, see [About Admin Groups](#).
3. Click the **Administration** → **Administrators** tab.
4. Select the group that you created and click the Edit icon. The out-of-the-box group for SAML authenticated users is **saml-group**.
5. Click the **SAML** tab.
6. Select **Auto Create User** for users in the IdP group to be automatically created in NIOS. When a new IdP user logs in to NIOS, that user is created in NIOS.
7. Select **Persist Auto Created User after logout** if you want to retain the SAML user accounts in NIOS even after the session times out. The session timeout value is specified in the **Session Timeout** field when you enabled SAML authentication. For more information, see [Enabling SAML Authentication](#). If you do not select the **Persist Auto Created User after logout** checkbox, if the session times out, users for whom the **SAML Only** option was selected in the **Authentication Type** field, are deleted from NIOS. Not selecting the **Persist Auto Created User after logout** checkbox also deletes the user account along with all the scheduled tasks associated with the user account when the user logs out of NIOS. For information about the Authentication Type field, see [Creating Local Admins](#).
8. Click **Save & Close**.



#### Note

If you select the **Persist Auto Created User after logout** checkbox and the session times out, you must manually verify whether the user account exists in IdP or not. If the user account is deleted from IdP, then you must manually delete the account in NIOS.

## Authenticating SAML Users

When you create administrators, you can authenticate them either as a SAML-only administrator or as a SAML/local administrator. Depending on the authentication type, administrators can log in using either the SSO Login button or the Login button. For more information see [Creating Local Admins](#).

## Administrative Permissions for Named ACLs

Only superusers and limited-access users with Read/Write permission to All Named ACLs and Read/Write permission to the corresponding objects and operations can manage named ACLs and their ACEs.

For information about access control and named ACLs, see [Configuring Access Control](#). The following table lists the operations and required permissions for managing named ACLs.

### Administration Permissions

Tasks	Grid Member(s)	All Named ACLs	DNS Views	Related DNS objects	File Distribution
Create, modify, and delete named ACLs for all supported operations	RW	RW	RW	RW	RW
View named ACLs and ACEs	RW	RO	RO	RO	RO

## Administrative Permissions for DNS Threat Protection

You can grant read-only or read/write permission, or deny access to the following resources:

- Grid Security Properties—Applies to the Grid and its members.
- Member Security Properties—Applies to the Grid members only.

For information about setting permissions, see [About Administrative Permissions](#). The following table lists the tasks admins can perform and the required permissions for the threat protection service.

### Permissions for hardware-based Threat Protection Service

Tasks	Grid Security Properties	Member Security Properties
View Grid security properties	RO	
Update Grid Security properties	RW	
View member security properties for specific Grid members	RO	RO
Update member security properties for specific Grid members	RW	RW



Tasks	Grid Security Properties	Member Security Properties
Start and stop threat protection service for a Grid member	RW	RW
Publish rules for a Grid member	RW	RW
View rule categories and rules for the Grid	RO	
Enable and disable rules for the Grid	RW	
Update rule versions for any rules on the Grid	RW	
Revert to a previous rule version for any rules on the Grid	RW	
Modify configuration parameters, such as action and severity, for rules on the Grid	RW	
Create custom rules from rule templates for the Grid	RW	
Delete custom rules for the Grid	RW	
View rule categories and rules on a Grid member	RO	RO
Enable and disable rules on a Grid member	RW	RW
Update rule versions for any rules on a Grid member	RW	RW
Revert to a previous rule version for any rules on a Grid member	RW	RW
Modify configuration parameters, such as action and severity, for rules on a Grid member	RW	RW
View threat protection related event statistics on a Grid member	RO	RO
Upgrade rulesets for a Grid	RW	

*Permissions for Software ADP*

Tasks	Grid Security Properties	Member Security Properties
View the list of Threat Protection profiles in the Profiles Viewer	RO	RO

Tasks	Grid Security Properties	Member Security Properties
View profile settings in the Threat Protection Profile Editor	RO	
Create a Threat Protection profile	RW	
Clone a Threat Protection profile from an existing profile (This also clones all settings for the ruleset from an old profile.)	RW	
Clone a Threat Protection profile from an existing member settings	RW	
Update the profile settings (name, comment, events per second, disable multiple TCP DNS request, list of members)	RW	
Change the ruleset that is assigned to a profile (This internally merges all customizations for an old ruleset to a new ruleset.)	RW	
View the profile rules and rule settings	RO	
Enable/disable rules in the profile	RW	
Change the rule parameters for rules in the profile (action, log severity, events per second etc.)	RW	
Merge two profiles	RW	
Assign/remove a profile from Member Security properties	RW	RW
Delete a profile	RW	

### Administrative Permissions for DNS Threat Analytics

Only superusers and limited-access users with Read/Write permission can manage Threat Analytics service. You can grant read-only or read/write permission, or deny access to the following:

- Grid Threat Analytics Properties—Applies to the Grid and its members.

For information about setting permissions, [Managing Permissions](#). The following table lists the tasks admins can perform and the required permissions for the threat analytics service.

#### Permissions for Threat Analytics Service

Tasks	Grid Threat Analytics Properties	RPZ Zones	Grid Members	DNS Views
View Grid Threat Analytics properties	RO		RO	

Tasks	Grid Threat Analytics Properties	RPZ Zones	Grid Members	DNS Views
Update Threat Analytics properties	RW	RW	RW	RW
Start and stop Threat Analytics service	RW		RW	
Create an RPZ and use it as mitigation blacklist feed	RW	RW	RW	RW
View whitelisted domains	RO		RO	
Move blacklisted domains to the whitelist	RW	RW		
Update Threat Analytics module and whitelist sets	RW			
Viewing threat analytics module and whitelist versions	RO			
Define the Threat Analytics Update policy	RW			
Manually Upload Threat Analytics Updates	RW			

### Administrative Permissions for All Rulesets

You can grant permissions for individual ruleset objects to admin users. NIOS provides a global permission ALL Rulesets for admin groups. To perform operations on an NXDOMAIN ruleset, a blacklist rule, or an RPZ ruleset, you must have permission to the rule or ruleset to which the ruleset object belongs.

### Administrative Permissions for Cloud Objects

You can grant read-only or read/write permission, or deny access to the following cloud related objects:

- All tenant—Applies to all tenants.
- Per tenant object—Applies to selected tenants only.

You need appropriate permissions to make changes to all tenants or a specific tenant object through Grid Manager. Note that following:

- Update the tenant object: You must have permission on all tenant objects or the specific tenant object that is being updated. When you modify any tenant object associated with a tenant, it requires explicit Read/Write permission on the specific tenant object, whether permission on the associated tenant exists or not.
- Read-Only: This operation returns all tenant objects if you have the all tenant permission, or returns only specific tenant objects for which you have access.
- Permission for a tenant object implicitly gives the read permission to any object that is associated with the tenant.
- Note that creating and deletion operations do not need any permission.

## Administrative Permissions for CSV Import

The following CSV import permissions are applicable to local users on the NIOS appliance. They are not valid for AD or RADIUS users.

- When you delete a user, CSV import tasks associated with the respective user are not deleted. Superusers can access these tasks. Pending CSV import tasks will not be executed due to authentication failure.
- If you change user permissions, the pending and running CSV import tasks are executed, but may finish with errors. Note that the appliance re-establishes database transaction after every 500 lines are imported. If you delete or modify the respective user entry between these transactions, the rest of the import may fail.
- Superusers can manage any stopped, failed, or completed CSV import tasks belonging to a deleted user.

## Administrative Permissions for Reporting

NIOS supports global and object permissions for reporting: Reports and Searches. Consider the following when applying permissions for reports and searches:

- When you view a DNS report, DHCP report, or search options, you can view all the data in the corresponding report or search results, even if you do not have permissions to view the DNS or DHCP objects.
- A limited-access user needs relevant permissions to view any report or search.
- When you grant a read-only permission for a specific report to a user, the user receives read-only access to everything the report displays.
- A user can view report data returned by a search operation even when the user is denied access to the search operation.
- You can only view pre-defined reports, global reports and customized reports or searches. A user, including superuser, cannot view customized reports or searches of other users.
- You cannot edit pre-defined reports.
- Only a superuser can edit global reports or searches.
- Only an owner of the customized report or search operation can modify the report.
- Superuser may create reports up to system wide limit.
- A limited-access admin can create up to five reports by default. You can configure this limitation through Grid Manager.

## Administrative Permissions for VLAN Management

A superuser has RW access for all VLAN related objects.

Global permissions and object permissions are supported for VLAN views, ranges, and objects. Permissions set at specific levels override global permissions for VLAN ranges, views, and objects. For more information see [About Administrative Permissions](#).

A user who belongs to the VLAN Admin role has the following permissions set:

Permission Type	Object	Task	Permission
VLAN Permissions	VLAN view	Creating VLAN views	RW
		Editing VLAN views	RW
		Deleting VLAN views	RW
VLAN Permissions	VLAN range	Creating VLAN ranges	RW
		Editing VLAN ranges	RW

Permission Type	Object	Task	Permission
		Deleting VLAN ranges	RW
VLAN Permissions	VLAN object	Creating VLAN objects	RW
		Editing VLAN objects	RW
		Deleting VLAN objects	RW
IPAM Permissions	Network view	Creating network views	RO
		Editing network views	RO
		Deleting network views	RO
IPAM Permissions	IPv4/IPv6 network		RW

For information about VLAN management, see [VLAN Management](#).

## Administrative Permissions for SAML

A user who belongs to the SAML Admin role has the following permissions set:

Permission Type	Resource Type	Permission
SAML Permissions	SAML Admin group	RW
SAML Permissions	SAML Authentication Services	RW

For information about SAML authentication, see [Authenticating Admins Using SAML](#).

## Deploying a Grid


To deploy a Grid, it is important to understand what a Grid is, how to create a Grid Master and add members, and how to manage the Grid. This section explains in detail about a grid and covers the topics:

- [Creating a Grid Master](#)
- [About Grids](#)
- [About HA Pairs](#)
- [Adding Grid Members](#)
- [Configuring an IPv6-only Grid](#)
- [Auto-Provisioning NIOS Appliances](#)
- [Pre-Provisioning NIOS and vNIOS Appliances](#)
- [Configuring a Grid](#)
- [Managing a Grid](#)
- [Grid Bandwidth Considerations](#)
- [Automatic Software Version Coordination](#)
- [NAT Groups](#)

## Creating a Grid Master

To create a Grid, you first create a Grid Master, and then add members. Although the Grid Master can be a single appliance (a "single master"), a more resilient design is to use an HA pair (an "HA master") to provide hardware redundancy. For information about HA pairs, see [About HA Pairs](#). The basic procedure for forming two appliances into an HA master is shown in the [Initially Configuring a Pair of Appliances as a Grid Master](#) figure. You can create a Grid Master in either IPv4, IPv6, or dual mode (IPv4 and IPv6). An IPv4 Grid Master uses IPv4 as the Grid communication protocol, so it supports IPv4 and dual mode Grid members. An IPv6 Grid Master uses IPv6 as the Grid communication protocol, so it supports IPv6 and dual mode Grid members. A dual mode Grid Master supports IPv4, IPv6, and dual mode Grid members. You can set either IPv4 or IPv6 as the communication protocol. All Infoblox hardware platforms, except for appliances with a 50 GB disk, support configuration as a Grid Master or Grid Master candidate. For information about which vNIOS appliance supports configuration as a Grid Master, see [vNIOS Appliances](#).

*Initially Configuring a Pair of Appliances as a Grid Master*



- 2 Connect Node 1 to the switch, log in to its default IP address (192.168.1.2), check that a Grid license is installed, and then configure the following:
- VIP address, netmask, gateway
  - Hostname
  - HA and LAN1 addresses of Node 1
  - HA and LAN1 addresses of Node 2
  - VRID (virtual router ID)
  - NTP settings
  - Grid name
  - Shared secret

3 After you configure Node 1, it listens for three seconds for VRRP advertisements containing its VRID number. When it does not receive any, it assumes the active role in the HA pair and starts sending advertisements.

Note:  
For more information about VRRP advertisements, see VRRP Advertisements page

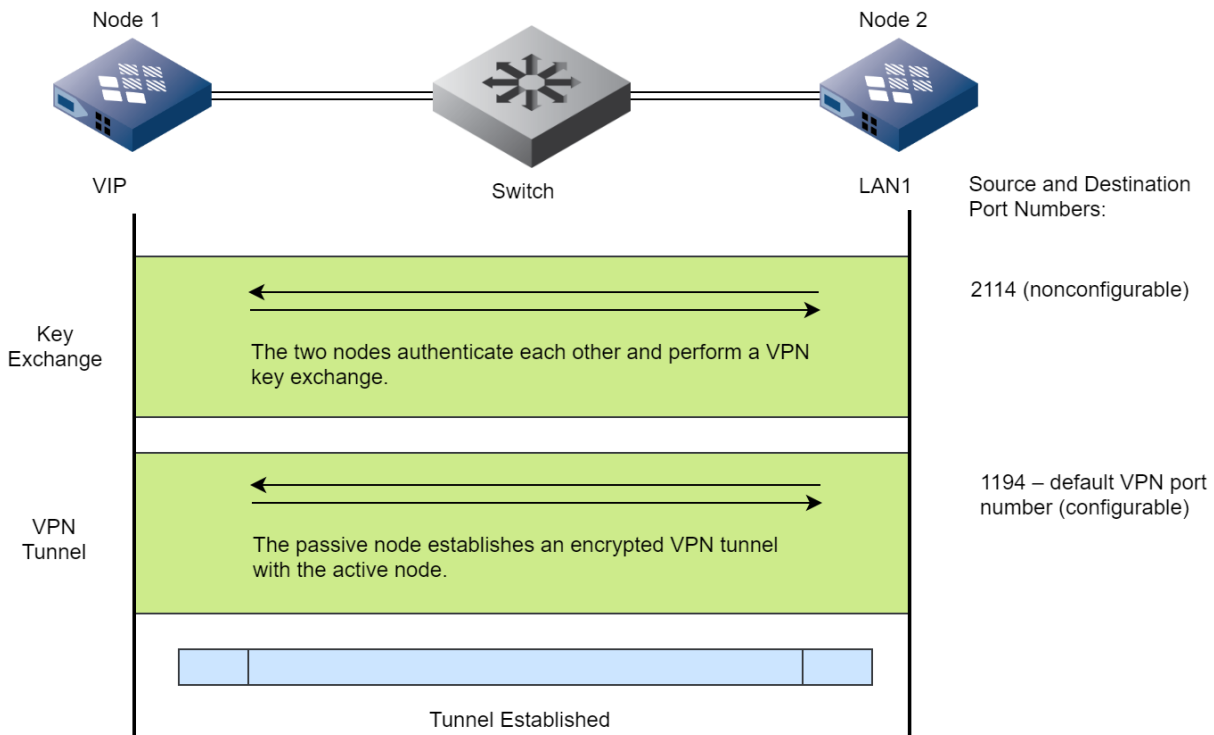
- 4 Connect Node 2 to the switch, log in to its default IP address (192.168.1.2), check that a Grid license is installed, and then configure the following:
- VIP address (for Node 1)
  - LAN1 address, netmask, gateway
  - Hostname
  - Grid name
  - Shared secret

Note:  
Because you do not set the VRID for Node 2, it cannot listen for VRRP advertisements yet. It learns its VRID after it joins the Grid and downloads the database from Node 1. Then, when Node 2 receives an advertisement containing its VRID from Node 1, it assumes the passive role in the HA pair.

5 After you configure Node 2, it contacts the VIP address on Node 1 and initiates a mutual authentication of the nodes using the shared secret. The nodes then construct an encrypted VPN tunnel to secure Grid communications.

After the two nodes form an HA pair, Node 2 initiates a key exchange and creates an encrypted VPN tunnel with Node 1. The two nodes communicate between the VIP interface linked to the HA port on Node 1 and the LAN1 port on Node 2. The initialization of VPN communications between the two nodes is shown in the below figure.

*Establishing a VPN Tunnel for Grid Communications*



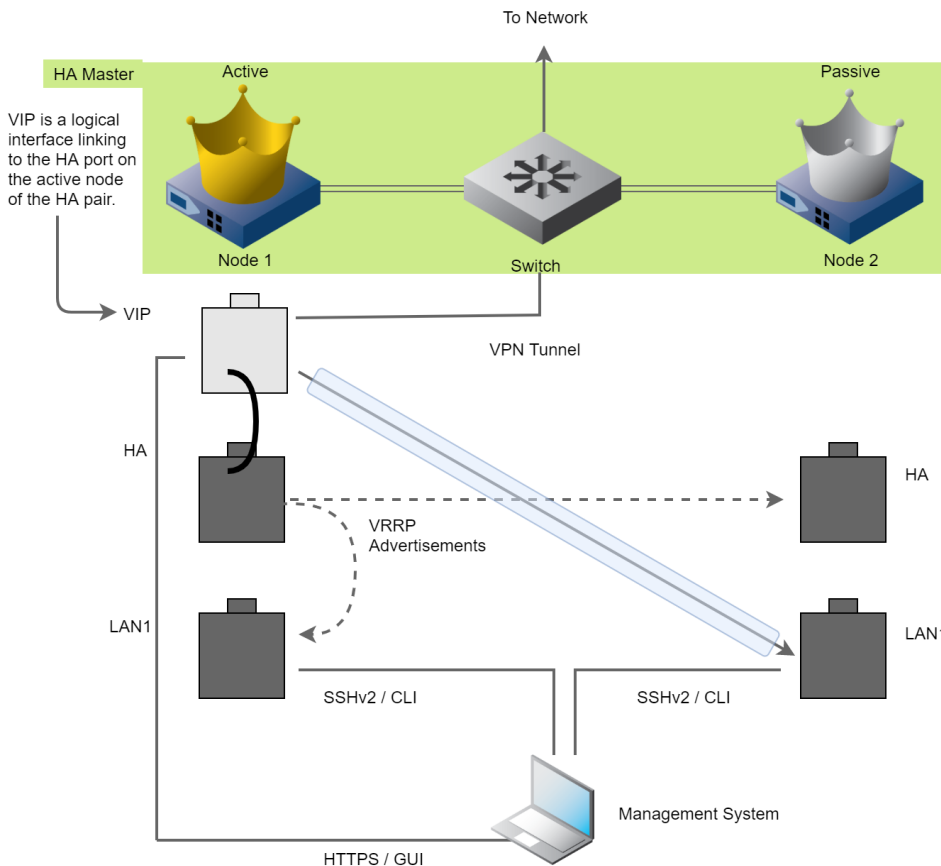
After the nodes establish a VPN tunnel between themselves, Node 1 sends Node 2 its entire database (its configuration settings and service data). Since the configuration contains the VRID (virtual router ID) for the HA pair, Node 2 starts listening for VRRP advertisements containing that VRID number. Also, since Node 1 is already sending such advertisements, Node 2 receives one and assumes the passive role in the HA pair.

After the initial transmission of its database, Node 1 continues to send Node 2 real-time database updates through the VPN tunnel.

Node 1 maintains the synchronization of the database throughout the Grid—which, at this point, has no other members—sends VRRP advertisements indicating its physical and network health, and—if configured to do so— provides network services. Node 2 maintains a state of readiness to assume the mastership in the event of a failover. You can see the flow of HA and Grid-related traffic from ports on the active node to ports on the passive node in the Traffic and Ports that an HA Grid Master Uses figure. This illustration also shows the ports that you can use for management traffic and network service.

*Traffic and Ports that an HA Grid Master Uses*





Note: If you enable the MGMT port, you can only make an HTTPS connection to the IP address of the active node. If you try to connect to the IP address of the passive node, the appliance redirects your browser to the IP address of the active node.

SSHv2, however, behaves differently from HTTPS. If you enable the MGMT port and define its network settings for both nodes in the HA pair, you can make an SSHv2 connection to the IP addresses of the LAN1 and MGMT ports on both the active and passive node

From the management system, you can manage the active node of the HA Master by making an HTTPS connection to the VIP interface and using the GUI, and by making an SSHv2 connection to the LAN1 port (and MGMT port, if enabled) and using the CLI. If you enable the MGMT port on an HA pair, you can make an HTTPS connection through the MGMT port on the active node, and you can make an SSHv2 connection through the LAN1 or MGMT port on the active and passive nodes.



#### Note

For information about enabling and using the MGMT port, SSH, and the Infoblox GUI, see [Using the MGMT Port, Restricting GUI/API Access](#), and [Logging on to the NIOS UI](#).

## Port Numbers for Grid Communication

If connectivity between Grid members must pass through a firewall, the firewall policies must allow the initial key exchange and subsequent VPN traffic to pass. The key exchange uses UDP with a source and destination port of 2114. VPN traffic uses UDP with a default source and destination port of 1194. The VPN port number is configurable.

To configure the VPN port number, complete the following:

1. On the **Grid** tab -> **Grid Manager** tab, expand the Toolbar, and then select **Grid Properties** -> **Edit**.
2. On the **General** tab of the *Grid Properties* editor, type a new port number in the **VPN Port** field.
3. Save the configuration.
4. When Grid Manager displays a warning indicating that a product restart is required, click **Yes** to continue. The product automatically restarts.

A member and master first perform a handshake to authenticate each other and exchange encryption keys. Then, they build an encrypted VPN tunnel between themselves. The member typically initiates both of these connections. The

master only initiates a key exchange if you manually promote a member to the role of master, (see [Promoting a Master Candidate](#)). The figure *Establishing a VPN Tunnel for Grid Communications* shows the typical connection exchange and default port usage not only between the two nodes forming an HA pair but also between a member and master when the member joins a Grid.

The member and master key exchange occurs when an appliance joins a Grid, during master promotion, and when a member reconnects to a Grid after becoming disconnected. At all other times, Grid-related communications occur through encrypted VPN tunnels.

## Grid Setup Wizard

The *Grid Setup Wizard* simplifies configuring a Grid. You can use it to configure an HA or single Grid Master and to join appliances to a Grid. The *Grid Setup Wizard* appears when you first log in to the appliance. After that, you can access it at any time as follows:

1. On the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Expand the Toolbar and select **Grid Properties** -> **Setup (Grid Setup Wizard)**.

## Creating an HA Grid Master

To create a Grid, you first create a Grid Master and then add members. Although you can define a single appliance as a Grid Master, using an HA pair provides hardware redundancy for this vital component of a Grid. You can create an HA Master in either IPv4, IPv6, or in dual mode. An IPv4 HA Master uses IPv4 as the communication protocol between the two nodes and for Grid communication. An IPv6 HA pair uses IPv6 as the communication protocol between the two nodes and for Grid communication. But in a dual mode HA pair, you can select either IPv4 or IPv6 as the communication protocol between the two nodes and for Grid communication. The following procedure explains how to put two NIOS appliances on the network and use the *Grid Setup Wizard* to configure them as Nodes 1 and 2 to form an HA Grid Master. For information about which vNIOS appliance supports configuration as an HA Grid Master, see [vNIOS Appliances](#).

## Configuring the Connecting Switch

To ensure that VRRP (Virtual Router Redundancy Protocol) works properly, configure the following settings at the port level for all the connecting switch ports (HA, LAN1, and LAN2):

- **Spanning Tree Protocol:** Disable this or enable this with Portfast. For vendor specific information, search for "HA" in the Infoblox Knowledge Base system at <https://support.infoblox.com>.
- **Trunking:** Disable.  
Note that if VLAN tagging is enabled on an Infoblox HA appliance, you must enable trunking at the port level.  
**EtherChannel:** Disable.
- **IGMP Snooping:** Disable.
- **DHCP Snooping:** Disable or Enable Trust Interface.  
Note that you must disable DHCP Snooping to successfully run DHCP services on the Grid. For more information about DHCP services, see [About Infoblox DHCP Services](#).
- **Port Channeling:** Disable.
- **Speed and Duplex settings:** Match these settings on both the Infoblox appliance and switch.
- Disable other dynamic and proprietary protocols that might interrupt the forwarding of packets.



### Note

By default, a NIOS appliance automatically negotiates the optimal connection speed and transmission type (full or half duplex) on the physical links between its LAN1 or LAN1 (VLAN), HA, and MGMT ports and the Ethernet ports on the connecting switch. If the two appliances fail to auto-negotiate the optimal settings, see [Modifying Ethernet Port Settings](#) for steps you can take to resolve the problem.

## Placing Both Appliances on the Network

1. Connect the power cable from each NIOS appliance to a power source and turn on the power. If possible, connect the appliances to separate power circuits. If one power circuit fails, the other might still be operative.
2. Connect Ethernet cables from the LAN1 port and the HA port on each appliance to a switch on the network.
3. Use the LCD on one appliance or make a console connection to it and configure the network settings of its LAN1 port so that it is on the local subnet and you can reach it on the network. LCD supports only IPv4 addressing and not IPv6 addressing. You can configure IPv6 address for the appliance through CLI or GUI. IPv4 addressing is supported on the LCD; ensure that you have the correct network address values before configuration of the appliance.  
Note that for details about using the LCD and console, refer to the installation guide that shipped with your product.
4. Similarly, configure the LAN1 port on the other appliance so that it is in the same subnet as the first appliance.
5. Connect your management system to the network so that it can reach the IP addresses of the LAN1 ports on both appliances.

## HA Master – Node 1

1. On your management system, open a browser window, and connect to `https://ip_addr`, where `ip_addr` is the IP address of the LAN1 port on Node 1. IPv4 and IPv6 values are valid, based on the LAN1 port configuration.
2. Log in using the default username and password: **admin** and **infoblox**. For detailed information about logging in to the GUI, see [Logging on to the NIOS UI](#).
3. Read the *Infoblox End-User License Agreement (EULA)*, and then click **I Accept**.  
Note that if you want to view the privacy policy of Infoblox, then on the EULA screen, click **Infoblox Privacy Policy**. Grid Manager displays the policy on a new browser tab.
4. Click **OK**. The *Grid Setup* wizard appears.
5. On the first screen, select **Configure a Grid Master** and click **Next**.
6. On the next screen, specify the Grid properties and click **Next**:
  - **Grid Name**: Enter a text string that the two appliances use to authenticate each other when establishing a VPN tunnel between them. The default Grid name is **Infoblox**.
  - **Shared Secret**: Enter a text string that both appliances use as a shared secret to authenticate each other when establishing a VPN tunnel between them. The default shared secret is **test**.
  - **Confirm Shared Secret**: Enter the shared secret again.
  - **Hostname**: Enter a valid domain name for the appliance.
  - **Type of Network Connectivity**: Select the type of network connectivity from the drop-down list:
    - **IPv4 and IPv6**: Select this to configure a dual mode HA Master.
    - **IPv4**: Select this to configure an IPv4 HA Master.
    - **IPv6**: Select this to configure an IPv6 HA Master.
  - **Is the Grid Master an HA pair?**: Select **Yes**.
    - **Send HA and Grid Communication over**: This field is displayed only when you are configuring a dual mode HA pair. Select either **IPv4** or **IPv6** as the communication protocol for VRRP advertisements.

### Notes

- Infoblox recommends that you back up the configuration after you convert a Grid to a different mode.
  - Restoring the old backup by performing a forced restore, may prevent the Grid members from rejoining the Grid Master after the restore.
7. On the next screen, specify the network properties and click **Next**:
    - **Virtual Router ID**: Enter the VRID (virtual router ID). This must be a unique VRID number—from 1 to 255—for this subnet.
    - **Ports and Addresses**: This table lists the network interfaces based on the type of network connectivity of the HA Master.  
For IPv4 HA Master, specify the network information for VIP (IPv4), Node1 HA (IPv4), Node2 HA (IPv4), Node1 LAN1 (IPv4), and Node2 LAN1 (IPv4) interfaces.  
For IPv6 HA Master, specify the network information for VIP (IPv6), Node1 LAN1 (IPv6), and Node2 LAN1 (IPv6) interfaces.  
For a dual mode HA Master, if you select **IPv4** in the **Send HA and Grid Communication over** field, specify the network information for the following interfaces: VIP (IPv4), Node1 HA (IPv4), Node1 LAN1 (IPv4),

Node2 HA (IPv4), Node2 LAN1 (IPv4), VIP (IPv6), Node1 LAN1 (IPv6), and Node2 LAN1 (IPv6) interfaces.

For a dual mode HA Master, if you select **IPv6** in the **Send HA and Grid Communication over** field, specify the network information for the following interfaces: VIP (IPv4), Node1 LAN1 (IPv4), Node2 LAN1 (IPv4), VIP (IPv6), Node1 LAN1 (IPv6), and Node2 LAN1 (IPv6) interfaces.

Enter correct information for the following by clicking the field:

- **Interface:** Displays the name of the interface. You cannot modify this.
- **Address:** Type the IPv4 or IPv6 address depending on the type of interface.
- **Subnet Mask (IPv4) or Prefix Length (IPv6):** Specify an appropriate subnet mask for IPv4 address or prefix length for IPv6 address. The prefix length ranges from 2 to 127.
- **Gateway:** Type the IPv4 or IPv6 address of the default gateway depending on the type of interface. For the IPv6 interface, you can also type **Automatic** to enable the appliance to acquire the IPv6 address of the default gateway and the link MTU from router advertisements.

#### Note

You can now define a link-local address as the default IPv6 gateway and isolate the LAN segment so that the local router can provide global addressing and access to the network and Internet. This is supported for both LAN1, LAN2, and VLAN interfaces, as well as LAN1, LAN2, VLAN in the failover mode. However, the link-local address does not support the following:

- IPv6 link local gateway for the MGMT interface.
  - IPv6 link local is not supported for addresses. It supported only for gateways.
- **VLAN Tag:** For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
  - **Port Settings:** From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
8. Optionally, enter a new password and click **Next**. The password must be a single string (no spaces) that is at least four characters long.
  9. Select the time zone of the Grid Master and indicate whether the Grid Master synchronizes its time with an NTP (Network Time Protocol) server.
    - If you choose to enable NTP, click the Add icon and enter the IP address of an NTP server. Entries may be an IPv4 or IPv6 address. You can enter IP addresses for multiple NTP servers.
    - If you choose to disable NTP, set the date and time for the appliance.
    - Click **Next**.
  10. The last screen displays the settings you specified in the previous panels of the wizard. Verify that the information is correct and click **Finish**. The application restarts after you click **Finish**.

Note that the *Grid Setup Wizard* provides options such as not changing the default password and manually entering the time and date. However, changing the password and using an NTP server improves security and accuracy (respectively), and so these choices are presented here.

Record and retain this information in a safe place. If you forget the shared secret, you need to contact Infoblox Technical Support for help. When you add an appliance to the Grid, you must configure it with the same Grid name, shared secret, and VPN port number that you configure on the Grid Master.
  11. Close the management window. The configuration for Node 1 is complete.

## HA Master – Node 2

1. On your management system, open a new browser window, and connect to [https://ip\\_addr](https://ip_addr), where *ip\_addr* is the IP address of the LAN1 port on Node 2. IPv4 or IPv6 values are valid.

When you enter an IPv6 address, enclose the address in square brackets (as in [https://\[2001:db8::256:ABCD:EF12:34:1\]](https://[2001:db8::256:ABCD:EF12:34:1])).
2. Log in using the default username and password **admin** and **infoblox**.
3. Read the Infoblox *End-User License Agreement* (EULA), and then click **I Accept**.

Note that if you want to view the privacy policy of Infoblox, then on the EULA screen, click **Infoblox Privacy Policy**. Grid Manager displays the policy on a new browser tab.
4. Click **OK**. The *Grid Setup* wizard appears.
5. On the first screen, select **Join Existing Grid** and click **Next**.

6. On the next screen, specify the Grid properties and click **Next**.
  - **Grid Name:** Enter a text string that the two appliances use to authenticate each other when establishing a VPN tunnel between them. This must match the Grid name you entered for Node 1.
  - **Grid Master's IP Address:** Enter the same VIP you entered for Node 1.
  - **Shared Secret:** Enter a text string that both appliances use as a shared secret to authenticate each other when establishing a VPN tunnel between them. This must match your entry in Node 1.
7. On the next screen verify the IP address settings of the member and click **Next**.  
The last screen displays the settings you specified in the previous panels of the wizard.
8. Verify that the information is correct and click **Finish**.  
The setup of the HA master is complete. From now on, when you make an HTTPS connection to the HA pair, use the VIP address.

The communication protocol for all the services in a dual mode (IPv4 and IPv6) HA Master is the same protocol as the one used for VRRP advertisements. For example, if you select **IPv4** in the **Send HA and Grid Communication over** field in step 2 of the *Grid Setup* wizard, then IPv4 is set as the communication protocol for all the services. However, you can override the communication protocol for all the services in a dual mode HA Master. For more information, see [Changing the Communication Protocol for a Dual Mode Appliance](#).

## Creating a Single Grid Master

Although using an HA master is ideal because of the hardware redundancy it provides, you can also use a single appliance as the Grid Master. You can create a single Grid Master in either IPv4, IPv6, or dual mode (IPv4 and IPv6). Infoblox recommends frequent backups if the Grid Master is a single appliance, and there is no Master Candidate. For information about which vNIOS appliance supports configuration as a single Grid Master, see [vNIOS Appliances](#).

Setting up an appliance as a single Grid Master is very easy. If the appliance has the DNSone package with the Grid upgrade, it is already a Grid Master. You simply need to define the network settings for its LAN1 port. The various procedures for defining the network settings for the LAN1 port of a single independent appliance apply here as well. Therefore, you can use any of the following procedures to define the network settings for the LAN1 port of the appliance that you want to make a single Grid Master:

- **LCD:** See [Method 1 – Using the LCD](#). (LCD configuration does not support IPv6 address entry.)
- **Console port:** See [Method 2 – Using the CLI](#).

You can also use the NIOS *Grid Setup Wizard* to create a single Grid Master. In addition to providing a simple method accompanied by helpful information, the setup wizard allows you to change the admin password and configure time settings for the appliance.

## Using the Setup Wizard

To create a single Grid Master using the *Grid Setup Wizard*, complete the following:

1. Connect the power cable from the NIOS appliance to a power source and turn on the power.
2. Connect an Ethernet cable from the LAN1 port on the appliance to a switch on the network.
3. If you have not changed the default IP address (192.168.1.2/24) of the LAN1 port through the LCD or CLI—and the subnet to which you connect the appliance does not happen to be 192.168.1.0/24—put your management system in the 192.168.1.0/24 subnet and connect an Ethernet cable between your management system and the NIOS appliance.
4. Open a web browser and make an HTTPS connection to the IP address of the LAN1 port. To reach the default IP address, enter: `*https://192.168.1.2*`.  
Several certificate warnings appear during the login process. This is normal because the preloaded certificate is self-signed (and, therefore, is not in the trusted certificate stores in your browser) and has the host name [www.infoblox.com](#), which does not match the destination IP address you entered in step 3. To stop the warning messages from occurring each time you log in to the GUI, you can generate a new self-signed certificate or import a third-party certificate with a common name that matches the FQDN (fully qualified domain name) of the appliance. For information about certificates, see [Managing Certificates](#).
5. Log in using the default username **admin** and password **infoblox**.
6. Read the Infoblox *End-User License Agreement*, and then click **I Accept**. Note that if you want to view the privacy policy of Infoblox, then click **Infoblox Privacy Policy**. Grid Manager displays the policy on a new browser tab.

7. Click **OK**. The *Grid Setup Wizard* appears.
8. On the first screen, select **Configure a Grid Master** and click **Next**.
9. On the next screen, specify the Grid properties and click **Next**:
  - **Grid Name**: Enter a text string that the Grid Master and appliances joining the Grid use to authenticate each other when establishing a VPN tunnel between them. The default Grid name is **Infoblox**.
  - **Shared Secret**: Enter a text string that the Grid Master and appliances joining the Grid use as a shared secret to authenticate each other when establishing a VPN tunnel between them. The default shared secret is **test**.
  - **Confirm Shared Secret**: Enter the shared secret again.
  - **Hostname**: Enter a valid domain name for the appliance.
  - **Type of Network Connectivity**: Select the type of network connectivity for the Grid Master from the drop-down list:
    - **IPv4 and IPv6**: Select this to configure a dual mode Grid Master.
    - **IPv4**: Select this to configure an IPv4-only Grid Master.
    - **IPv6**: Select this to configure an IPv6-only Grid Master.
  - Note:
    - Infoblox recommends that you back up the configuration after you convert a Grid to a different mode.
    - Restoring the old backup by performing a forced restore, may prevent the Grid members from rejoining the Grid Master after the restore.
  - **Is the Grid Master an HA pair?**: Select **No**.
10. On the next screen, configure the network settings and click **Next**:
  - a. **Ports and Addresses**: This table lists the network interfaces based on the type of network connectivity of the Grid Master. For IPv4 Grid Master, specify the network information for LAN1 (IPv4) port and for IPv6 Grid Master, specify the network information for LAN1 (IPv6) port. For a dual mode Grid Master, specify the network information for both LAN1 (IPv4) and LAN1 (IPv6).  
Enter correct information for the following by clicking the field:
    - **Interface**: Displays the name of the interface. You cannot modify this.
    - **Address**: Type the IPv4 or IPv6 address depending on the type of interface. An IPv6 address is a 128-bit number in colon hexadecimal notation. It consists of eight 16-bit groups of hexadecimal digits separated by colons (example: 2001:db8:0000:0123:4567:89ab:0000:cdef or 2001:db8::123:4567:89ab:0:cdef).
    - **Subnet Mask (IPv4) or Prefix Length (IPv6)**: Specify an appropriate subnet mask for the IPv4 address or prefix length for the IPv6 address. The prefix length ranges from 2 to 127.
    - **Gateway**: Type the IPv4 or IPv6 address of the default gateway depending on the type of interface. For the IPv6 interface, you can also type **Automatic** to enable the appliance to acquire the IPv6 address of the default gateway and the link MTU from router advertisements. Note that you can now define a link-local address as the default IPv6 gateway and isolate the LAN segment, so the local router can provide global addressing and access to the network and Internet. For information about the link-local address limitations, see the **Gateway** field description in the *HA Master – Node 1* section.
    - **VLAN Tag**: For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
    - **Port Settings**: From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
11. Optionally, enter a new password and click **Next**. The password must be a single alphanumeric string (no spaces) that is at least four characters long. The password is case-sensitive.
12. Select the time zone of the Grid Master and indicate whether the Grid Master synchronizes its time with an NTP (Network Time Protocol) server, and then click **Next**.
  - If you choose to enable NTP, click the Add icon and enter the IP address of an NTP server. Entries may be an IPv4 or IPv6 address. You can enter IP addresses for multiple NTP servers.
  - If you choose to disable NTP, set the date and time for the appliance.
13. The last screen displays the settings you specified in the previous panels of the wizard. Verify that the information is correct and click **Finish**. The application restarts after you click **Finish**.



**Note**

The *Grid Setup* wizard provides options such as not changing the default password and manually entering the time and date. However, changing the password and using an NTP server improves security and accuracy (respectively), and so these choices are presented here.

Record and retain this information in a safe place. If you forget the shared secret, you need to contact Infoblox Technical Support for help. When you add an appliance to the Grid, you must configure it with the same Grid name, shared secret, and VPN port number that you configure on the Grid Master.

The last screen of the setup wizard states that the changed settings require the appliance to restart. When you click **Finish**, the appliance restarts.

The setup of the single master is complete. From now on, when you make an HTTPS connection to the appliance, use its new IP address.

In a dual mode Grid Master, the communication protocol for all the services is set to IPv4, by default. You can change the default communication protocol for the services. For information, see [Changing the Communication Protocol for a Dual Mode Appliance](#).

## About Grids

A Grid is a group of two or more NIOS appliances that share sections of a common, distributed, built-in database and which you configure and monitor through a single, secure point of access: the Grid Master. A Grid can include Infoblox appliances and vNIOS appliances. A vNIOS appliance is a non-Infoblox hardware platform running the vNIOS software package. For supported vNIOS platforms, see [vNIOS Appliances](#).

Infoblox appliances support both IPv4 and IPv6 networks and you can configure a Grid in one of the following modes:

- **IPv4-only:** An IPv4-only Grid uses IPv4 as the Grid communication protocol and it includes an IPv4 Grid Master and the Grid members, which can be either IPv4 or dual mode (IPv4 and IPv6) independent and HA appliances. Note that when you add a dual mode HA member to an IPv4-only Grid, the communication protocol between the two nodes of an HA pair must be IPv4.
- **IPv6-only:** An IPv6-only Grid uses IPv6 as the Grid communication protocol and it includes an IPv6 Grid Master and the Grid members, which can be either IPv6 or dual mode (IPv4 and IPv6) independent and HA appliances. Note that when you add a dual mode HA member to an IPv6-only Grid, the communication protocol between the two nodes of an HA pair must be IPv6.
- **IPv4 and IPv6 (Dual mode):** A dual mode Grid can use either IPv4 or IPv6 as the Grid communication protocol. A dual mode Grid includes a dual mode Grid Master and the Grid members, which can be either IPv4, IPv6, or dual mode independent and HA appliances.

**Note**

Infoblox appliances support IPv4 and IPv6 networking configurations in most deployments cited in this chapter. You can set the LAN1 port to an IPv6 address and use that address to access Grid Manager. All HA (high availability) operations can be applied across IPv6. Topics in this and following chapters generally use IPv4 examples. Also note that the LAN2, MGMT, and VLAN ports also support IPv6. DNS services are fully supported in IPv6 for the LAN1, LAN2, MGMT, and VLAN ports. DHCP services are fully supported in IPv6 for the LAN1 and LAN2 ports. Examples throughout this chapter use IPv4 addressing. Interfaces on NIOS appliances support both IPv4 and IPv6 transports and intra-Grid communication is based on the type of IP address used by the Grid member to join the Grid.

The following table summarizes the possible setups of a Grid configuration:

### *Possible Setups of Grid configuration*

Grid Configuration	VRRP Protocol for HA Pair	Grid Communication Protocol	Grid Connection via MGMT	Additional IPv4 Addresses	Additional IPv6 Addresses
IPv4 Grid Master	IPv4	IPv4	NA	Yes	Yes
IPv6 Grid Master	IPv6	IPv6	NA	Yes	Yes
Dual mode Grid Master	IPv4 or IPv6	IPv4 or IPv6	NA	Yes	Yes
IPv4 Grid member	IPv4	IPv4	IPv4	Yes	Yes
IPv6 Grid member	IPv6	IPv6	IPv6	Yes	Yes
Dual mode Grid member	IPv4 or IPv6	IPv4 or IPv6	IPv4 or IPv6	Yes	Yes



**Note**

Infoblox recommends that appliances with disk sizes below 250 GB must not be configured as Grid Masters.

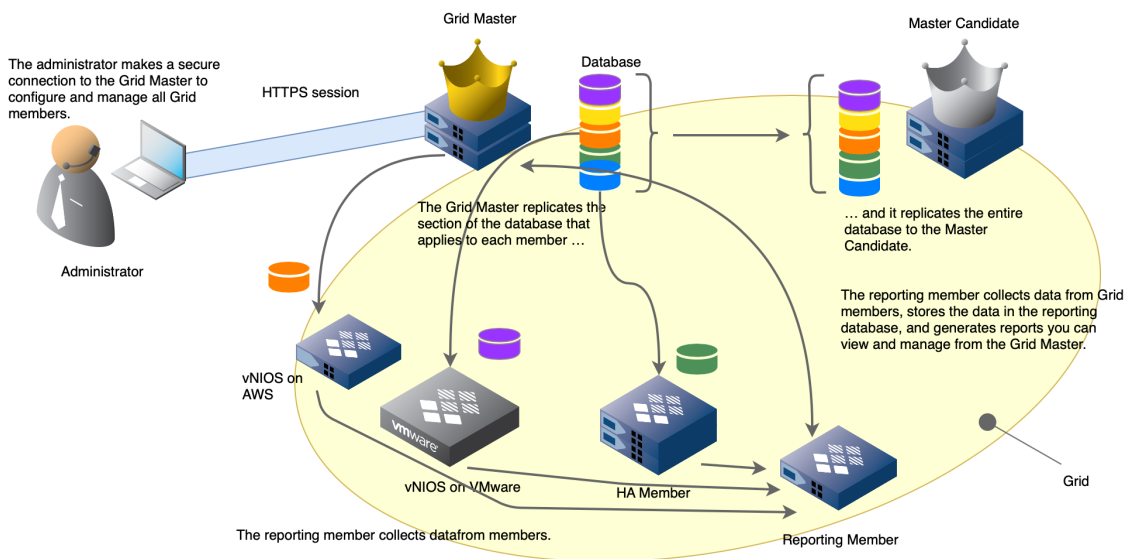
You can also add supported Reporting platforms as a logging and reporting devices in your Grid. Infoblox provides a few Infoblox platforms that you can use as the logging and reporting device. For information about the supported appliances, see [Configuring Reporting Clustering](#). Infoblox reporting solution supports both IPv4 and IPv6 networks and you can configure a reporting member in either IPv4, IPv6, or in dual mode (IPv4 and IPv6) network environment. An IPv4-only Grid uses IPv4 as the Grid communication protocol, so you can add an IPv4 or dual mode reporting member to an IPv4-only Grid. An IPv6-only Grid uses IPv6 as the Grid communication protocol, so you can add an IPv6 or dual mode reporting member to an IPv6-only Grid. However, a dual mode Grid can use either IPv4 or IPv6 as the Grid communication protocol, so you can add an IPv4, IPv6, or a dual mode reporting member to a dual mode Grid. The reporting appliance collects data from members in the Grid and stores the data in the database. It then uses the data to generate predefined and user-defined reports that you can access through Grid Manager. These reports provide useful information about the IPAM, DNS, DHCP, and system activities and usage in your Grid. For more information about reporting, see [Infoblox Reporting and Analytics](#).

Instead of manually provisioning IP addresses and DNS name spaces for network devices and interfaces, you can add Cloud Platform Appliances to leverage DNS and DHCP features of the Grid to manage your CMPs (Cloud Management Platforms). For information about the Infoblox Cloud Network Automation solution and supported Grid configurations, see [Deploying Cloud Network Automation](#).

The following figure shows the basic concept of a Grid, database distribution (or "replication"), and reporting.

*Grid and Partitioned Database Replication*



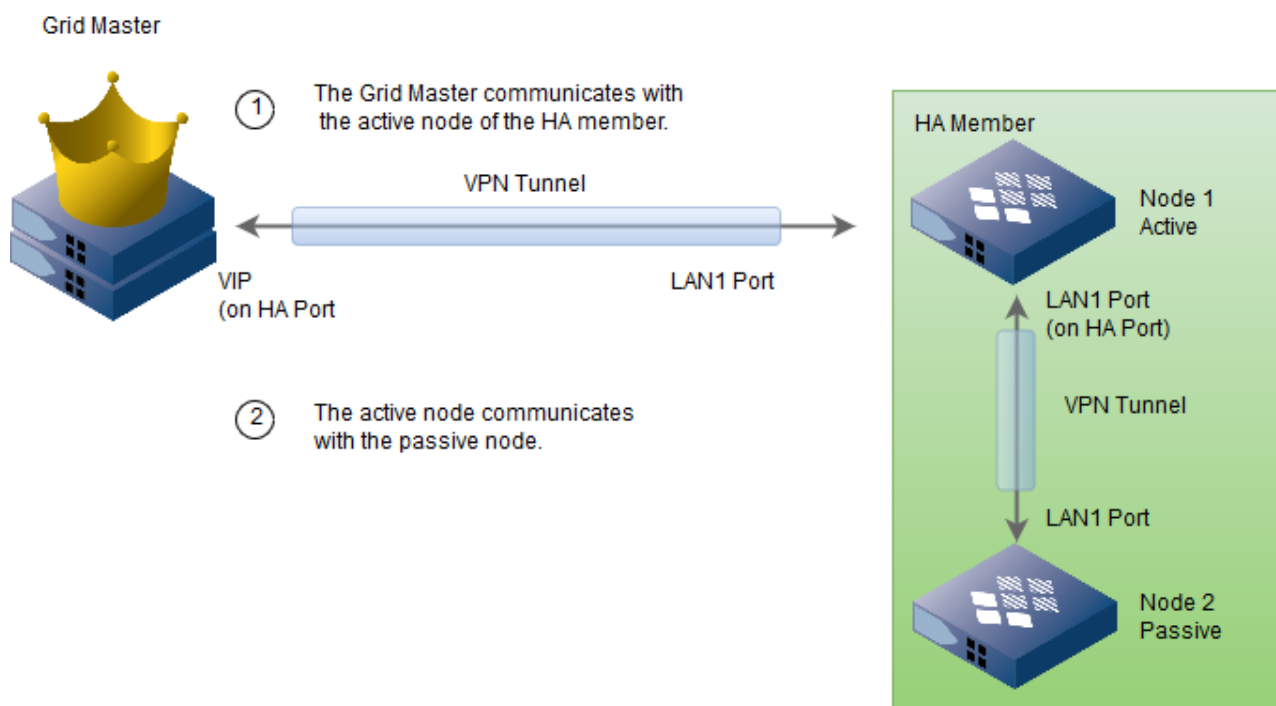


**Note:** In addition to the HTTPS session securing administrative traffic to the Grid Master, all Grid communications between the Grid Master and Grid members pass through encrypted VPN tunnels (not shown).

The Grid Master can be either an HA master or a single master; that is, an HA (high availability) pair or a single appliance. Similarly, a Grid member can be either a single member or an HA member. You can add single appliances and HA pairs to a Grid, forming single members and HA members respectively. A single Grid member can be either an Infoblox appliance or a vNIOS appliance. An HA Grid member can be a pair of Infoblox appliances or vNIOS appliances. For information, see [vNIOS Appliances](#).

The Grid Master communicates with every Grid member in a hub-and-spoke configuration. Intra-Grid communication is based on the type of IP address used by the Grid member to join the Grid Master. An IPv4-only Grid Master uses IPv4 and an IPv6-only Grid Master uses IPv6 for intra-Grid communication. However, a dual mode Grid Master uses either IPv4 or IPv6 depending on the IP address type used by the Grid member to join the Grid Master. For an HA member, the Grid Master communicates with the active node, which in turn communicates with the passive node, as shown in the following figure.

#### *Grid Communications to an HA Member*



When adding vNIOS appliances to a Grid, you centralize the management of core network services of the virtual appliances through the Grid Master. vNIOS appliances support most of the features of the Infoblox NIOS software, with some limitations as described in [vNIOS Appliances](#).

By default, Grid communications use the UDP transport with a source and destination port of 1194. This port number is configurable. For a port change to take effect, one of the following must occur: the HA master fails over, the single master reboots, or the Grid restarts services.

After adding an appliance or HA pair to a Grid, you no longer access the Infoblox GUI on that appliance. Instead, you access the GUI running on the Grid Master. Although you can create multiple administrator accounts to manage different services on various Grid members, all administrative access is through the Grid Master. So even if someone has administrative privileges to a single Grid member, that administrator must access the GUI running on the Grid Master to manage that member.

You can access the Infoblox GUI through an HTTPS connection to one of the following IP addresses and ports on the Grid Master:

- The VIP address, which links to the HA port on the active node of an HA Grid Master
- The IP address of the LAN1 port on a single Grid Master
- The IP address of the MGMT port (if enabled) of the active node of an HA or single Grid Master. See [Using the MGMT Port](#).

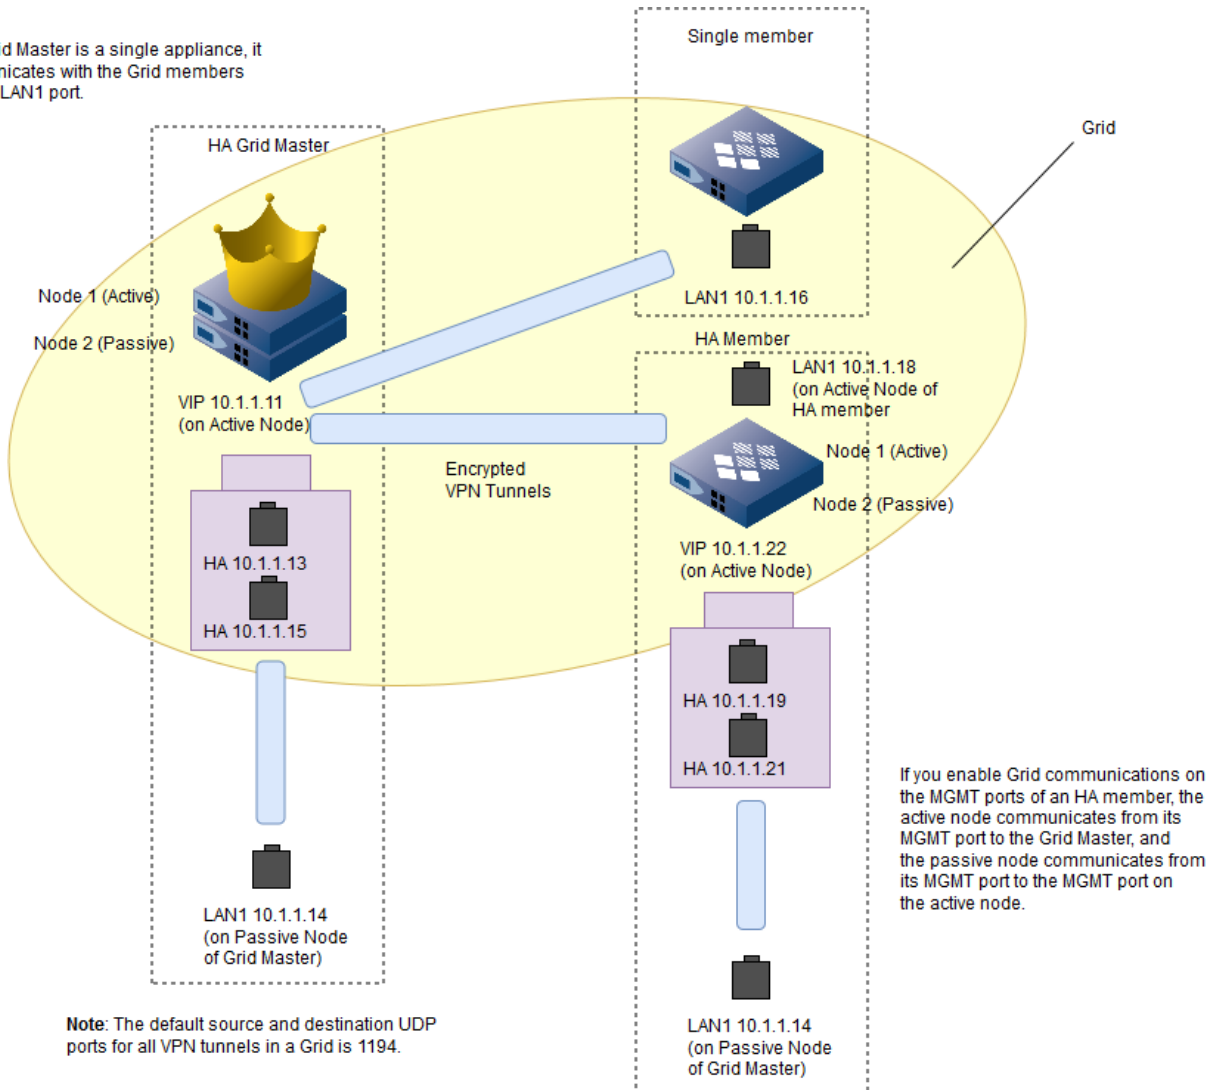
## Grid Communications

The Grid Master synchronizes data among all Grid members through encrypted VPN tunnels. The default source and destination UDP port number for VPN tunnels is 1194. You can continue using the default port number or change it. For example, if you have multiple Grids, you might want each Grid to use a different port so that you can set different firewall rules for each. Whatever port number you choose to use for the VPN tunnels in a Grid, all the tunnels in that Grid use that single port number.

Before an appliance or HA pair forms a tunnel with the master, they first authenticate each other using the Challenge-Response Authentication Mechanism (CRAM). The source and destination port number for this traffic is 2114. During the CRAM handshake, the master tells the appliance or HA pair what port number to use when building the subsequent VPN tunnel.

## VPN Tunnels within a Grid

If the Grid Master is a single appliance, it communicates with the Grid members from its LAN1 port.



Another type of traffic, which flows outside the tunnels, is the VRRP (Virtual Router Redundancy Protocol) advertisements that pass between the active and passive nodes in an HA pair. The VRRP advertisements act like heartbeats that convey the status of each node in an HA pair. If the active node fails, the passive node becomes active. The VIP (virtual IP) address for that pair then shifts from the previously active node to the currently active node.

## Master Grids

A Master Grid provides centralized management of multiple Grids. When a Grid is managed by a Master Grid, the Master Grid icon appears on the left side of the top panel of Multi-Grid Manager. Assuming you have permission, you can click this icon to access Multi-Grid Manager. In addition, the Toolbar provides several functions for joining the Master Grid, editing its properties and leaving the Master Grid. For more information about the Master Grid and these functions, refer to the *Multi-Grid Manager Administrator Guide*.

## About HA Pairs

You can configure two appliances as an HA (high availability) pair to provide hardware redundancy for core network services and Infoblox Advanced DNS Protection. For more information, see [About Infoblox Advanced DNS](#)

**Protection.** An HA pair can be a Grid Master, a Grid Master candidate, a Grid member, or an independent appliance. An HA pair can also comprise a physical appliance and a virtual appliance, two physical appliances, or two virtual appliances. The two nodes that form an HA pair—identified as Node 1 and Node 2—are in an active/passive configuration. The active node receives, processes, and responds to all service requests. The passive node constantly keeps its database synchronized with that of the active node, so it can take over services if a failover occurs. A failover is the reversal of the active/passive roles of each node; that is, when a failover occurs, the previously active node becomes passive and the previously passive node becomes active. You can configure an HA pair in either IPv4, IPv6, or in dual mode. An IPv4 HA pair uses IPv4 as the communication protocol between the two nodes and an IPv6 HA pair uses IPv6 as the communication protocol between the two nodes. But in a dual mode HA pair, you can select either IPv4 or IPv6 as the communication protocol between the two nodes. Note that when you add a dual mode HA member to a Grid, the communication protocol between the two nodes of an HA pair must be the same as the Grid communication protocol.



#### Note

HA Grid Master and HA Grid Master Candidate configurations are not supported when Threat Protection licenses are installed on the appliance.

When you configure an HA pair using the IB-4030-10GE (Rev-1 or Rev-2) appliance for DNS cache acceleration, the passive node does not operate with a pre-loaded cache or hot cache during a failover; it builds up the DNS cache over time. For more information about HA and other limitations for the IB-4030-10GE appliances, refer to the *Infoblox DNS Cache Acceleration Application Guide*. For Infoblox, only the active node in an HA pair handles DNS traffic. The passive node is in a standby mode ready to take over if a failover occurs.

The appliance uses the following components in the HA functionality:

- **bloxSYNC:** An Infoblox proprietary mechanism for secure, real-time synchronization of the database that maintains the data, system configuration, and protocol service configuration between the two nodes. With bloxSYNC, the nodes continuously synchronize changes of their configurations and states. When a failover occurs, the passive node can quickly take over services. For information, see [About HA Failover](#) below.
- **VRRP (Virtual Router Redundancy Protocol):** An industry-standard, MAC-level HA failover mechanism. VRRP utilizes the concept of an active and passive node that share a single VIP (virtual IP) address. When the active node that owns the VIP becomes unavailable, the passive node takes over the VIP and provides network core services. For information about VRRP, refer to *RFC3768, Virtual Router Redundancy Protocol (VRRP)* and see [VRRP Advertisements](#) below.

Using bloxSYNC and VRRP combined, if the active node fails or is taken offline for maintenance purposes, the passive node assumes the VIP and continues to respond to requests and services with minimal interruption. You can deploy an HA pair as a Grid Master, a Grid member, or an independent HA. To deploy an independent HA pair, see [Deploying an Independent HA Pair](#). To deploy an HA Grid Master, see [Creating a Grid Master](#).

## Planning for an HA Pair

To achieve high availability, the HA and LAN1 (or VLAN) ports on both the active and passive nodes are connected to switches on the same network or VLAN. Both nodes in an HA pair share a single VIP address and a virtual MAC address so they can appear as a single entity on the network. You can also assign IPv6 addresses for each of the active and passive nodes, in addition to the IPv6 VIP address.



#### Note

- Infoblox uses VRRP advertisements for the active and passive HA design. Therefore, all HA pairs must be located in the same location connected to the highly available switching infrastructure. Any other deployment is not supported without a written agreement with Infoblox. Contact Infoblox Technical Support for more information about other deployment support.
- You can enable ARP on the passive node of an HA pair and monitor its status externally. To enable ARP on the passive node of an HA pair, see [Enabling ARP on the Passive Node of an HA Pair](#) below.

In HA, each node must configure two addresses: the VRRP public address on the LAN1 interface and the VRRP HA address on the HA interface. An HA pair consists of a set of five IP addresses, all of which must belong to the same subnet. Each device in an HA pair joins the multicast address on both the HA and public interfaces.

As illustrated in the following figure, the VIP and virtual MAC addresses link to the HA port on each node. Select five IP addresses on the same network before you configure an HA pair, as follows:

- VIP: For core network services and for management purposes when the MGMT port is disabled. Both nodes share the same VIP. The VIP is the true public address in which services and daemons are active.
- Node 1 HA (active): Source IP for the VIP and VRRP advertisements. Listens on both its LAN and HA ports. For an active HA node, both the LAN interface/address and the HA interface/address belong to the VRRP multicast group.
- Node 1 LAN1 (active): For management through SSHv2 and listens for VRRP advertisements from the HA port.
- Node 2 HA (passive): Listens for VRRP advertisements on the LAN port. For a passive HA node, only the LAN interface/address belongs to the VRRP multicast group (using the LAN port's MAC address).
- Node 2 LAN1 (passive): Source IP for SSL VPN to the VIP of the active node and receives bloxSYNC from the VIP.




**Note**

An HA member connecting to the GRID Master over the Management port should have the LAN1 or HA ports connected to different physical switches to make sure the VRRP packets are exchanged correctly between the active and passive nodes. If the LAN1 or HA ports are connected to the same physical switch, you must configure the LAN1/LAN2 bonding to exchange the VRRP packets between the active and passive nodes.

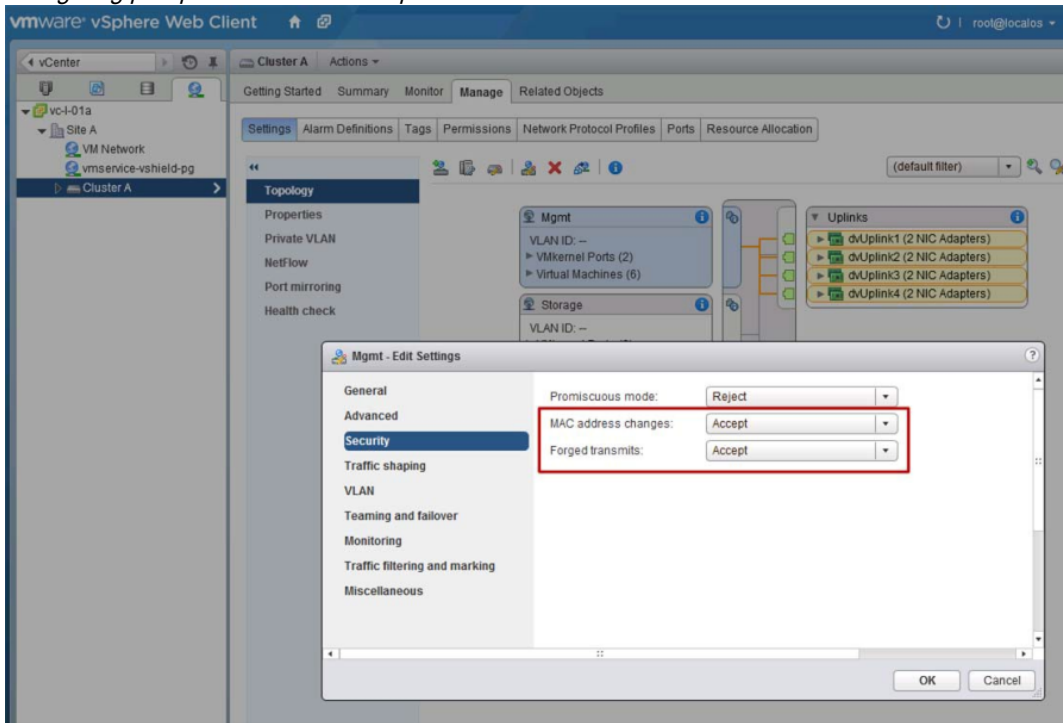
The above configuration holds good only for IPv4 VRRP configurations. IPv6 VRRP configurations require only three addresses: the VIP and the LAN1 interfaces. For the IPv6 dedicated HA interfaces, NIOS uses the link local IPv6 address which you do not need to provide.

*HA Pair*



When you deploy a vNIOS HA pair, ensure that the port connection allows for more than one MAC address per vNIC. For example, if you deploy a vNIOS HA pair in VMware vSphere, the port-profile to which the vNIOS HA and LAN ports connect should allow for more than one MAC address per vNIC. You can do this by changing the security settings of the port-group to accept "MAC address changes" and "Forged transmits," as illustrated in the following figure.

## Configuring port-profile in VMware vSphere



### Limitations of Using a Combination of a Physical Appliance and a Virtual Appliance for HA

Although you can use a combination of a physical and virtual appliance for an HA pair, using the deployment has the following limitations:

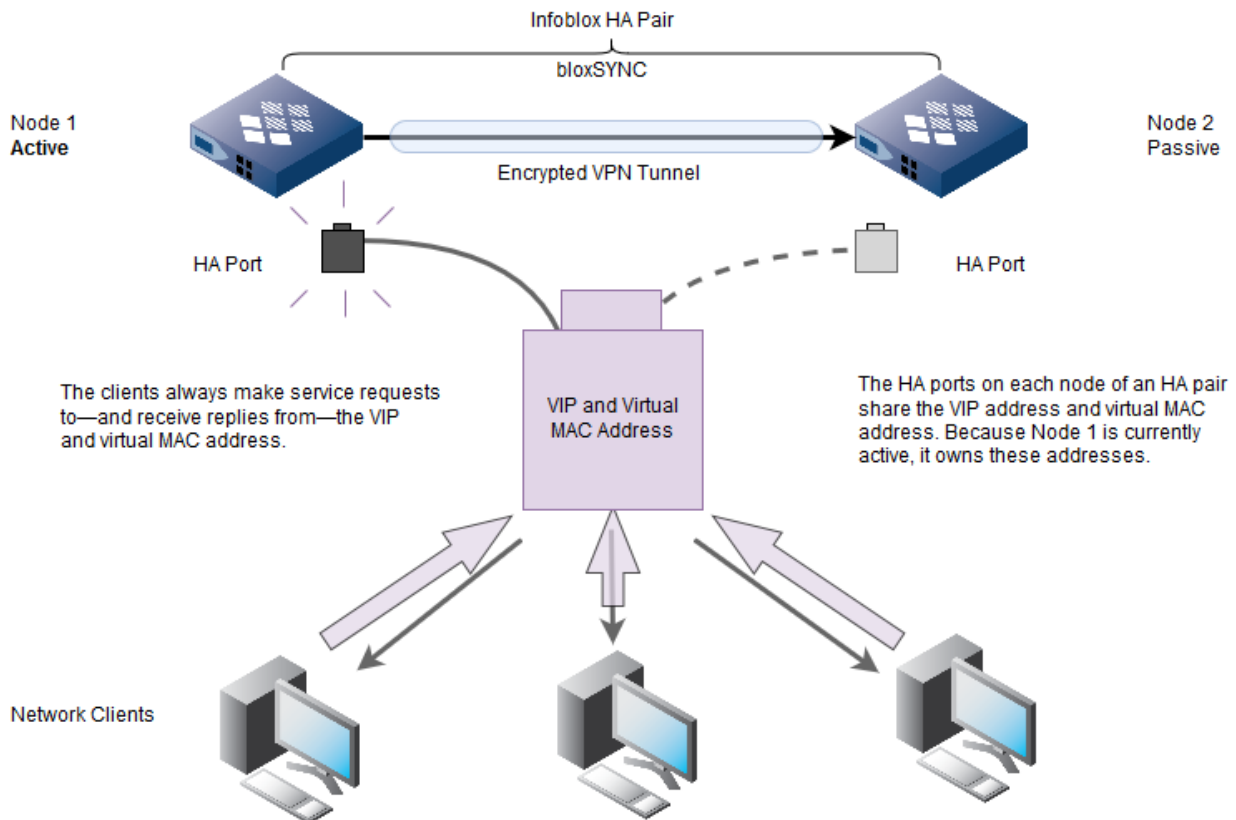
- Check whether the NIOS licenses that you subscribed to support both physical and virtual appliances.
- Ensure that the same licenses are installed on both the physical appliance and the virtual appliance.
- You cannot have tagged and untagged interfaces on the same subnet on VMware ESXi hypervisors.
- Virtual appliances do not support tagging.
- LOM (Lights Out Management) is not supported in a hybrid HA setup.
- DSCP (Differentiated Services Code Point) services are not supported on virtual appliances. Therefore, you cannot configure the DSCP value in an HA setup.
- Because port settings are not available for virtual appliances, you cannot join a node if the port settings are overridden.
- You cannot combine a platform on which Advanced DNS Protection hardware is running with a platform on which Advanced DNS Protection Software is running.
- You cannot configure MTU (Maximum Transmission Unit) in a hybrid HA setup.
- You cannot have a combination of an IB-FLEX and a non IB-FLEX appliance.
- Auto-provisioning is not supported on virtual appliances; therefore, you cannot use the auto-provisioning feature in a hybrid HA setup.
- A hybrid HA setup may cause some performance impact because hybrid HA performance depends on many factors such as the hardware on which the VM is running, the number of VMs contending for the same CPU, RAM, input/output resources, and the overhead generated by the virtualisation layer.
- Minor performance differences are expected between the two nodes of a hybrid HA pair. Hybrid HA performance may vary, and it depends on the hardware components on which different virtualization platforms are running and the performance delivered by Infoblox hardware appliances. Different use cases will produce different numbers (slightly increased or decreased CPU usage, disk access time, and so on). Such performance variation is expected and is not a cause of concern.

## About HA Failover


The appliance supports HA through bloxHA™, which provides a robust failover mechanism. As described in Planning for an HA Pair, both nodes in an HA pair share a single VIP address and a virtual MAC address. The node that is currently active is the one whose HA port owns the VIP address and virtual MAC address. When a failover occurs, these addresses shift from the HA port of the previous active node to the HA port of the new active node, as illustrated in the figure below.

**Note**  
For a vNIOS HA pair, you must configure both LAN1 and HA interfaces to operate. When there is a notification about failure in any one of the port, make sure that both of these ports are working. If one of the port is down and another port is still working, the HA pair believes its peer is active. But, there will be connectivity issues as one of the port is down. An HA failover occurs on vNIOS appliances only when both of these ports are down. For details about configuring these virtual NICs, refer to the *Infoblox Installation Guide vNIOS for VMware*.

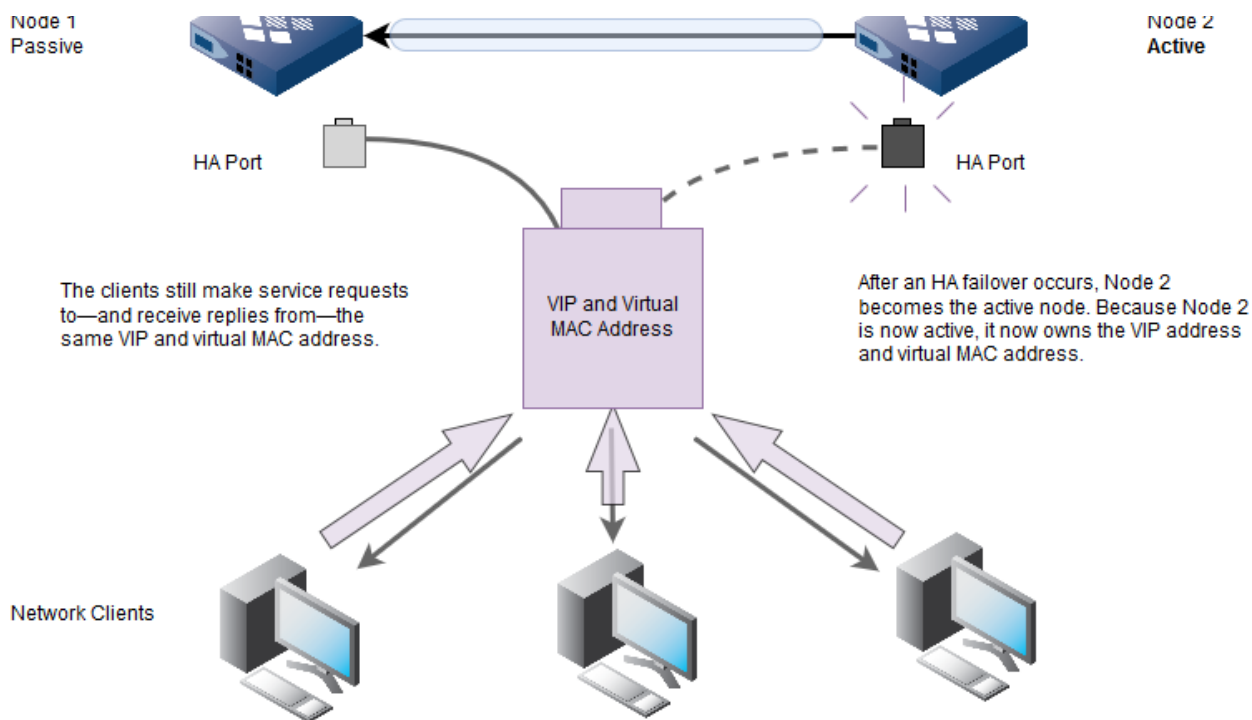
### VIP Address and Virtual MAC Address and HA Failover



### After an HA Failover







#### Enabling ARP on the Passive Node of an HA Pair


You can enable ARP (Address Resolution Protocol) on the passive node of an HA pair and monitor its status externally. For example, when the active node of an HA pair fails over to the passive node, you can ping the passive node from an external location and monitor its status. By default, ARP is disabled on the passive node of an HA pair. ARP settings on an HA member are preserved during a system restart or reboot, HA switch over, and upgrade. In addition, you do not need to restart the appliance when you modify ARP settings. When the active node becomes passive during an HA failover, ARP on an HA member inherits the settings configured in the database.

You can view detailed status for both nodes of an HA pair through the *Detailed Status* panel. To view the *Detailed Status* panel, from the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox -> Detailed Status icon in the horizontal navigation bar. In the *Detailed Status* panel, you can view ARP connectivity status for the passive node of an HA pair (Green = The passive HA node is connected to the local router; Yellow = The passive HA node fails to connect to the local router; Gray = ARP is disabled on the passive node of an HA pair). The passive HA node uses arping to test the ARP connectivity with the local router. If the local router is not configured, you may see false warnings even if the ARP connectivity is fine. In case of an ARP connectivity failure, the appliance sends an SNMP trap and an email notification, if configured.

Note that the ARP setting is not preserved on a passive HA node when you reset the appliance using the CLI command `reset all` or reset the database using the CLI command `reset database`.

To enable ARP on an HA passive node:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Select an HA member and click the Edit icon.
3. In the *Grid Member Properties* editor, select the **Network** tab -> **Advanced** tab and complete the following:
  - **Enable ARP on HA Passive Node?:** Select one of the following:
    - **Disable (default):** Select this to disable ARP on an HA passive node. This is selected by default.
    - **Enable (not recommended):** Select this to enable ARP on an HA passive node.
4. Save the configuration and click **Restart** if it appears at the top of the screen.



Note



For the Grid having an HA Grid Master and the **Enable ARP on HA Passive Node?** option enabled, if you try to restore from the HA Grid Master to a single node Grid Master the Grid breaks the configuration and if you try to recover the configuration the system becomes unusable. However, you can recover the configuration by resetting the database.



#### Warning

Enabling ARP on the passive node of an HA interface might affect VRRP on the local network and could cause the firewall to send false alerts.

## HA failover on DNS Nameservers

When an HA failover occurs on NIOS, there is an approximate 4-5 second time interval in which the network is adjusted for the new active node and the new passive node. During this failover period, the active node becomes unresponsive. After the new active node comes up on the network, the DNS service loads all Response Policy Zone (RPZ) files if RPZ is configured. The larger the RPZ files, the longer it takes to load them, and the longer it takes the DNS service to start serving DNS. For example, on a TE-1425 with RPZs that contain 15 million resource records, it can take approximately one and a half minutes to start serving DNS.

If your nameserver uses Grid replication to keep internal zones up to date and is not configured to use RPZ, then the delay before the DNS service starts serving DNS is slightly longer than it is for the HA failover itself.

## VRRP Advertisements

VRRP advertisements are periodic announcements of the availability of the HA node linked to the VIP. The two nodes in an HA pair include a VRID (virtual router ID) in all VRRP advertisements and use it to recognize VRRP advertisements intended for themselves. Only another appliance on the same subnet configured to use the same VRID responds to the announcements. The active node in an HA pair sends advertisements as multicast datagrams every second. It sends them from its HA port using the source IP address of the HA port (not from the VIP address) and the source MAC address 00:00:5e:00:01:vrrp\_id. The last two hexadecimal numbers in the source MAC address indicate the VRID number for this HA pair. For example, if the VRID number is 143, then the source MAC address is 00:00:5e:00:01:8f (8f in hexadecimal notation = 143 in decimal notation).

The destination MAC and IP addresses for all VRRP advertisements are 00:00:5e:00:01:12 and 224.0.0.18 (00:00:5e:00:02:12 and FF02::12 for IPv6 only configurations). Because a VRRP advertisement is a multicast datagram that can only be sent within the immediate logical broadcast domain, the nodes in an HA pair must be in the same subnet together.

As illustrated in the figure below, when you configure an HA pair, only the appliance configured to listen for VRRP advertisements with the same VRID number processes the datagrams, while all other appliances ignore them. The passive node in an Infoblox HA pair listens for these on its HA port and the active node listens on its LAN1 or LAN1 (VLAN) port. If the passive node does not receive three consecutive advertisements or if it receives an advertisement with the priority set to 0 (which occurs when you manually perform a forced failover or request the active node to restart, reboot, or shut down), it changes to the active state and assumes ownership of the VIP address and virtual MAC address.

If both nodes go offline, the one that comes online first becomes the active node. If they come online simultaneously, or if they enter a dual-active state—that is, a condition arises in which both appliances assume an active role and send VRRP advertisements, possibly because of network issues—then the appliance with the numerically higher VRRP priority becomes the active node. The priority is based on system status and events.


If both nodes have the same priority, then the appliance whose HA port has a numerically higher IP address becomes the active node. For example, if the IP address of the HA port on Node 1 is 10.1.1.80 and the IP address of the HA port on Node 2 is 10.1.1.20, then Node 1 becomes the active node.

For more information about VRRP, see *RFC 3768, Virtual Router Redundancy Protocol (VRRP)*.

### *VRRP Advertisements with a Unique VRID*

After you finish configuring Node 1 of the HA pair to use VRID 143—a number that is unique for this subnet—it starts listening for VRRP advertisements with that VRID. When it does not receive any for three seconds, it becomes the active node in the HA pair and begins multicasting VRRP advertisements with a VRID 143 from its HA port.

Any device on that subnet that is not configured to listen for VRRP advertisements with VRID 143 drops the packet.



After you finish configuring Node 2 to join the HA pair, it initiates a connection with Node 1. The two appliances establish a VPN tunnel between themselves, using the HA connection name and shared secret to authenticate each other. Node 2 downloads the database from Node 1 and learns its VRID. Node 2 then begins listening for VRRP advertisements on its HA port. When it receives an advertisement from Node 1, Node 2 recognizes it and becomes the passive node.



#### Note

For a dual mode (IPv4 and IPv6) HA Master or HA member, you can set either IPv4 or IPv6 for VRRP advertisements. The Grid members may restart all the services after the Grid Master has recovered from the dual-active state.

## Adding Grid Members

#### Note



You may provision a Port Reservation for the new Grid Member. When doing so, you select the device to which you expect the new Grid Member to connect; In the context of a Grid member, this device type is usually an Ethernet Switch or Switch-Router. The Add Grid Member Wizard provides a step in which you define the port reservation settings, as described in the following section Adding a Single Member. The process also can be applied when defining an HA pair, as described in the sections [Creating an HA Grid Master](#) and [Adding an HA Member](#) see below.

You can add single appliances and HA pairs to a Grid, forming single members and HA members, respectively. A single Grid member can be either an Infoblox appliance or a vNIOS appliance. You can configure Grid members in either IPv4, IPv6, or dual mode (IPv4 and IPv6). For information about which vNIOS appliance supports configuration as an HA Grid member, see [vNIOS Appliances](#).

You can also define an HA member on the Grid Master and then add two individual NIOS appliances to the Grid as Node 1 and Node 2 to complete the HA member you defined on the master.

New members inherit all settings that you create at the Grid level unless you override them at the member level. You can also define port reservations for the network infrastructure devices to which the Grid members will connect.

The process for adding either a single appliance or HA pair to a Grid involves the following steps:

1. Adding and configuring Grid members on the Grid Master. In addition to defining the network and appliance settings for a member, you can also configure service settings before you join the member or HA pair to the Grid.

2. Reserving a port on a switch or switch-router for connectivity to the Grid member.
3. Joining the appliance or HA pair to the Grid. This includes defining the VIP or IP address of the Grid Master, the Grid name, and the shared secret on the single appliance or HA pair. If an appliance or HA pair cannot join the Grid because of MTU (maximum transmission unit) limitations on its network link, you can reduce the MTU that the master uses when communicating with it. See [Setting the MTU for VPN Tunnels](#). If the Grid Master is behind a NAT device and there are members on both sides of that NAT device, you must create a NAT group, as described in [NAT Groups](#).

In a large-scale deployment of Grids across multiple sites, consider remotely provisioning your Grid members before joining them to the Grid. For more information about this feature, see [Auto-Provisioning NIOS Appliances](#).

In situations where you want to define certain configurations on an offline Grid member and associate DNS and DHCP data to the member before deploying it, you can use the pre-provisioning feature to accomplish this. For more information, see [Pre-Provisioning NIOS and vNIOS Appliances](#).

## Adding a Single Member

The basic steps necessary to add a single member are as follows:

1. Define the network settings of the LAN1 port of the single appliance on the Grid Master.
2. Initiate the join Grid operation during which you specify the VIP or IP address of the Grid Master, the Grid name, and the shared secret on the single appliance. For information, see [Joining Appliances to the Grid](#) below.

On the Grid Master, you can configure any service settings such as DNS zones and records, DHCP networks and address ranges, and other services for a member before or after you join the appliance to the Grid. The basic steps for adding a single member are presented in the following section.

For information on how to configure a vNIOS appliance as a Grid member, refer to the [Quick Start Guide for Installing vNIOS Software on VMware Platforms](#).

## Configuring a Single Member on the Grid Master

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Expand the Toolbar and click **Add** -> **Add Grid Member**.
3. In the *Add Grid Member* wizard, enter the following and click **Next**:
  - **Member Type**: Specify the appliance type of the Grid member. If the member is an Infoblox appliance, select **Infoblox**, which is the default. For a vNIOS appliance, select **Virtual NIOS** for vNIOS virtual appliances, including cloud virtual appliances.
  - **Host Name**: Type the FQDN (fully qualified domain name) of the appliance that you are adding to the Grid. Ensure to use a host name containing characters (A-Z or a-z), numbers (0–9), ".", "-", and "\_". Additional special characters cannot be used.
  - **Time Zone**: If the Grid member is in a different time zone from the Grid, click **Override** and select a time zone.
  - **Comment**: Type a comment that provides some useful information about the appliance, such as its location.
  - **Master Candidate**: Select this option to designate this appliance as a Master Candidate. For supported vNIOS appliances, see [vNIOS Appliances](#).
4. Enter the following information about the member that you are adding to the Grid and click **Next**:
  - **Type of Network Connectivity**: Select the type of network connectivity for the Grid member from the drop-down list:
    - **IPv4 and IPv6**: Select this to configure a dual mode Grid member.
    - **IPv4**: Select this to configure an IPv4 Grid member.
    - **IPv6**: Select this to configure an IPv6 Grid member.

Note that Infoblox recommends that you back up the configuration after you convert a Grid to a different mode. Restoring the old backup by performing a forced restore, may prevent the Grid members from rejoining the Grid Master after the restore.

- **Standalone Member**: Select this option.
- **Required Ports and Addresses**: This table lists the network interfaces based on the type of network connectivity of the Grid member. For IPv4 Grid member, specify the network information for LAN1 (IPv4) port and for IPv6

Grid member, specify the network information for LAN1 (IPv6) port. For a dual mode Grid member, specify the network information for both LAN1 (IPv4), and LAN1 (IPv6).

Enter correct information for the following by clicking the field:

- **Interface:** Displays the name of the interface. You cannot modify this.
- **Address:** Type the IPv4 or IPv6 address depending on the type of interface. An IPv6 address is a 128-bit number in colon hexadecimal notation. It consists of eight 16-bit groups of hexadecimal digits separated by colons (example: 2001:db8:0000:0123:4567:89ab:0000:cdef or 2001:db8::123:4567:89ab:0:cdef).
- **Subnet Mask (IPv4) or Prefix Length (IPv6):** Specify an appropriate subnet mask for IPv4 address or prefix length for IPv6 address. The prefix length ranges from 2 to 127.
- **Gateway:** Type the IPv4 or IPv6 address of the default gateway depending on the type of interface. For IPv6 interface, you can also type **Automatic** to enable the appliance to acquire the IPv6 address of the default gateway and the link MTU from router advertisements.
- **VLAN Tag:** For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
- **Port Settings:** From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
- **DSCP Value:** Displays the Grid DSCP value, if configured. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP, see [Configuring Ethernet Ports](#).

5. In the Port Reservation page, do the following:

Begin by checking the Reserve Port checkbox. Note that reserving a switch port does not guarantee its availability once the device must connect. The port is automatically assigned for connectivity to the LAN1 port on the appliance.

Optionally, you can skip connecting port configuration by clicking **Next**. Click the **Clear** button to remove the selected device from the configuration.

- Click the **Select Device** button to choose the device for which the port reservation will be associated. You should know the identity of the device to which the Infoblox appliance will connect before taking this step. For Grid member connectivity, the chosen device should be either a switch or a switch-router.
- After choosing the device, choose the **Interface** with which the reservation will be bound. The drop-down list shows only interfaces that are most recently found to be available by Grid Manager during the last Discovery cycle. This list will not include any ports that are Administratively Up and Operationally Up or that are otherwise already assigned to other networks or Objects.
- The Wizard page also shows a list of any VLANs that are currently configured in the chosen device (**The following VLANs are configured**). This Wizard page does not allow the definition of new VLANs for port configuration—only the assignment of an existing VLAN in the device to your new port reservation. (Recall that you may specify the **VLAN Tag** across which Grid member traffic will travel, when you specified the Grid member information in Step 2 of the Wizard.)
- Check the **Configure Port** checkbox to define specific Port Control settings for the port reservation.
- Choose the **Data VLAN** and/or the **Voice VLAN** settings you may need for the port assignment. Depending on the selected device, the **Voice VLAN** field may or may not appear.
- Set the **Admin Status** to **Up** if you need to activate the port after assignment in the current task.
  - All Port Control operations require CLI credentials to be entered into Grid Manager. Because some IPAM and DHCP Objects will use Port Control features as part of object creation, CLI credentials are automatically leveraged as part of discovery. Ensure you have the correct sets of CLI credentials for devices in your network.
- Enter a **Description** for the port assignment. Infoblox recommends doing so to help other technicians to recognize the port assignment event.
- When finished, click **Next** to continue in the wizard.

6. Optionally, define extensible attributes. For information, see [About Extensible Attributes](#).

7. The final step for adding a Grid member is to define when the associated Port Configuration task executes. You may execute it immediately or schedule it for another time and date.

- To create the new port configuration immediately, select **Now**. The port control task is automatically synchronized to take place at the same time as the activation of the new Grid member.
- You can choose to have Grid Manager execute the port control task at a later time. To do so, select **Later**. Choose a Selected time by entering or selecting a Start Date (click the calendar icon to choose a calendar date) and a Start Time, and choose a Time Zone.

8. Choose one of the following from the **Save &...** drop-down button menu:

- Click **Save & Close** to add the single member to the Grid and close the wizard (this is the default).
- Click **Save & Edit** to add the single member to the Grid and launch the editor. You can configure additional properties, such as the MTU size, or add the member to a NAT group.
- Click **Save & New** to add the single member to the Grid and launch the wizard again to add another member.

The communication protocol for all the services in a dual mode (IPv4 and IPv6) Grid member is set to IPv4, by default. You can change the default communication protocol for all the services. For information, see [Communication Protocol for a Dual Mode Appliance](#) below.

## Adding an HA Member

The basic steps necessary to add an HA member are as follows:

1. Define the network settings of the HA pair on the Grid Master.
2. Initiate the join Grid operation, during which you specify the VIP or IP address of the Grid Master, the Grid name, and the shared secret on the HA pair. For information, see [Joining Appliances to the Grid](#) below.

In addition, on the Grid Master you can configure the service settings such as DNS zones and records, DHCP networks and address ranges, and so on for a member before or after you join the HA pair to the Grid. The basic steps for adding an HA member are presented below.



### Note

The procedure for adding an HA pair to a Grid when it uses the MGMT port of the active node for Grid communications differs slightly from that described below. See [Grid Communications](#).

## Configuring an HA Member on the Grid Master

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Expand the Toolbar and click **Add** -> **Add Grid Member**.
3. In the *Add Grid Member* wizard, enter the following and click **Next**:
  - **Member Type**: Specify the appliance type of the Grid member. If the member is an Infoblox appliance, select Infoblox, which is the default. For a vNIOS appliance on VMware, select **Virtual NIOS**.
  - **Host Name**: Type the FQDN (fully qualified domain name) for the HA member. Ensure to use a host name containing characters (A-Z or a-z), numbers (0–9), ".", "-", and "\_". Additional special characters cannot be used.
  - **Time Zone**: If you want the Grid member to have a different time zone, click **Override** and select a time zone.
  - **Comment**: Type a comment that provides some useful information about the appliance, such as its location.
  - **Master Candidate**: select this checkbox to designate this appliance as a Master Candidate. For supported vNIOS appliances, see [vNIOS Appliances](#).
4. Enter the following information about the member that you are adding to the Grid and click **Next**:
  - **Type of Network Connectivity**: Select the type of network connectivity for the HA member from the drop-down list:
    - **IPv4 and IPv6**: Select this to configure a dual mode HA member.
    - **IPv4**: Select this to configure an IPv4 HA member.
    - **IPv6**: Select this to configure an IPv6 HA member.
  - **High Availability Pair**: Select this option.
    - **Virtual Router ID**: Enter a unique VRID number—from 1 to 255—for the local subnet.

- **Send HA and Grid Communication Over:** This field is displayed only when you are configuring a dual mode HA member. Select either **IPv4** or **IPv6** as the communication protocol for VRRP advertisements and for joining the Grid Master.

Note that Infoblox recommends that you back up the configuration after you convert a Grid to a different mode.

Restoring the old backup by performing a forced restore, may prevent the Grid members from rejoining the Grid Master after the restore.

- **Required Ports and Addresses:** This table lists the network interfaces based on the type of network connectivity. For IPv4 HA member, specify the network information for VIP (IPv4), Node1 HA (IPv4), Node2 HA (IPv4), Node1 LAN1 (IPv4), and Node2 LAN1 (IPv4) interfaces. For IPv6 HA member, specify the network information for VIP (IPv6), Node1 LAN1 (IPv6), and Node2 LAN1 (IPv6) interfaces.  
For a dual mode HA member, if you select **IPv4** in the **Send HA and Grid Communication over** field, specify the network information for the following interfaces: VIP (IPv4), Node1 HA (IPv4), Node1 LAN1 (IPv4), Node2 HA (IPv4), Node2 LAN1 (IPv4), VIP (IPv6), Node1 LAN1 (IPv6), and Node2 LAN1 (IPv6) interfaces.  
For a dual mode HA member, if you select **IPv6** in the **Send HA and Grid Communication over** field, specify the network information for the following interfaces: VIP (IPv4), Node1 LAN1 (IPv4), Node2 LAN1 (IPv4), VIP (IPv6), Node1 LAN1 (IPv6), and Node2 LAN1 (IPv6) ports.  
Enter correct information for the following by clicking the field:
  - **Interface:** Displays the name of the interface. You cannot modify this.
  - **Address:** Type the IPv4 or IPv6 address depending on the type of interface. An IPv6 address is a 128-bit number in colon hexadecimal notation. It consists of eight 16-bit groups of hexadecimal digits separated by colons (example: 2001:db8:0000:0123:4567:89ab:0000:cdef or 2001:db8::123:4567:89ab:0:cdef).
  - **Subnet Mask (IPv4) or Prefix Length (IPv6):** Specify an appropriate subnet mask for IPv4 interface or prefix length for IPv6 interface. The prefix length ranges from 2 to 127.
  - **Gateway:** Type the IPv4 or IPv6 address of the default gateway depending on the type of interface. For IPv6 interface, you can also type **Automatic** to enable the appliance to acquire the IPv6 address of the default gateway and the link MTU from router advertisements.
  - **VLAN Tag:** For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
  - **Port Settings:** From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
  - **DSCP Value:** Displays the Grid DSCP value, if configured. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP, see [Implementing Quality of Service Using DSCP](#).  
Note that when the system operates in HA mode, should the IPv6-addressed VIP value be deleted, the IPv6 address of the HA port will also be deleted.

5. Optionally, define extensible attributes. For information, see [Using Extensible Attributes](#).

6. Do one of the following:

- Click **Save & Edit** to add the HA member to the Grid and launch the editor. You can configure additional properties, such as the MTU size, or add the member to a NAT group.
- Click **Save & New** to add the HA member to the Grid and launch the wizard again to add another member.
- Click **Save & Close** to add the HA member to the Grid and close the wizard.

The communication protocol for all the services in a dual mode (IPv4 and IPv6) HA member is the same protocol as the one that is used for VRRP advertisements. For example, if you select **IPv4** in the **Send HA and Grid Communication over** field in step 2 of the *Add Grid Member* wizard, then IPv4 is set as the communication protocol for all the services. However, you can override the communication protocol for all the services in a dual mode HA member. For information, see [Communication Protocol for a Dual Mode Appliance](#) below.



## Changing the Member Type

When you change the **Member Type** from *Infoblox* to *Virtual NIOS*, Infoblox displays an error indicating that the network port of a vNIOS member must be set to Automatic. If you encounter this error, follow the steps mentioned below to change the **Member Type** to **Virtual NIOS**:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Expand the **Toolbar** and click **Add** -> **Add Grid Member**.
3. In the *Add Grid Member* wizard, leave the **Member Type** as **Infoblox**, fill other details and click **Next**.
4. In the **Network** tab select **High Availability Pair**.
5. Change the port settings to **Automatic** for **Node1 HA**.
6. Select **Standalone Member**.
7. Click **Previous** and change the **Member Type** to **vNIOS**.

## Changing the Communication Protocol for a Dual Mode Appliance

You can change the default communication protocol for a dual mode appliance. You can force the appliance to use a specific protocol to join the Grid Master and for the reporting services. But for services with two types of resolution (A and AAAA records), you can set the preferred communication protocol.

To change the communication protocol for a dual mode appliance:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox -> Edit icon.
2. In the *Grid Member Properties* editor, select the **Network** tab -> **Basic** tab, and then complete the following:
  - **Communication Protocol Settings and Preferences:** This setting is not applicable for an HA pair. Select either **IPv4** or **IPv6** from the drop-down list. This setting will force the appliance to use the specified protocol for Grid and reporting services and this is the preferred protocol for services with two types of resolution (A and AAAA records).
  - **Customized Settings:** Select this and do the following:
    - **Always use this Communications Protocol for:** For a Grid Master, you can select either **IPv4** or **IPv6** from the **Reporting** drop-down list. This setting will force the Grid Master to use the specified communication protocol for reporting service. For a Grid member, you can select either **IPv4** or **IPv6** from the **Grid** and **Reporting** drop-down list. This setting will force the Grid member to use the specified communication protocol for Grid and reporting service.
    - **Always Prefer this Communications Protocol for:** This field lists the services which has two types of resolution (A and AAAA records). Select either **IPv4** or **IPv6** from the drop-down list for the service which you want the appliance to use this as the preferred communication protocol. The appliance uses the preferred protocol first for the service.

## Joining Appliances to the Grid

Grid members can join the Grid using IPv4 protocol in an IPv4-only Grid and using IPv6 protocol in an IPv6-only Grid. In a dual mode Grid, the Grid members may join the Grid using IPv4 or IPv6. Similarly, a Grid Master candidate can join the Grid using IPv4 in an IPv4-only Grid and using IPv6 in an IPv6-only Grid. But for a Grid Master candidate to join a dual mode Grid, it should be configured in dual mode. If you have configured the MGMT port for the Grid member, then the Grid member can join the Grid using the MGMT port. You can use the Grid Setup Wizard or access the *Join Grid* dialog box to join appliances to a Grid. The Grid Setup Wizard launches when you first log in to an appliance. You can also launch it from the Toolbar as described in [Grid Setup Wizard](#).

To join a single appliance and HA pair to a Grid using the Grid Manager GUI:

1. Log in to the appliance or HA pair that you want to add to the Grid. The appliance or HA pair must be online and able to reach the Grid Master.
2. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
3. Expand the Toolbar and click **Join Grid**.
4. In the *Join Grid* dialog box, enter the following:
  - **Virtual IP of Grid Master:** Type the VIP address of the HA Grid Master or the LAN1 address of the single Grid Master for the Grid to which you want to add the appliance. Entries may be an IPv4 or IPv6 address.
  - **Grid Name:** Type the name of the Grid.
  - **Grid Shared Secret:** Type the shared secret of the Grid.

- **Use MGMT port to join Grid:** If you have already enabled the MGMT port (see [Grid Communications](#)), this option becomes available. Select it to connect to the Grid through the MGMT port.
5. Click **OK** to begin the join operation.  
To confirm that the appliance has successfully joined the Grid, log in to the Grid Master and navigate to the **Grid** tab, select the **Grid Manager** -> **Members** tab. This panel lists the Grid members. Check the icon in the Status column of the newly added member. (green = the appliance has joined the Grid and is functioning properly; yellow = the appliance is in the process of joining the Grid; red = the appliance has not joined the Grid). You can also use the CLI command set network to join an appliance to a Grid.

To join a single appliance and HA pair to a Grid using the Grid Setup Wizard:

1. Log in to the appliance or HA pair that you want to add to the Grid. The appliance or HA pair must be online and able to reach the Grid Master.
2. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
3. Expand the Toolbar and click **Grid Properties** -> **Setup (Grid Setup Wizard)**.
4. On the next screen, specify the Grid properties and click **Next**
  - **Grid Name:** Enter a text string that the two appliances use to authenticate each other when establishing a VPN tunnel between them. This must match the Grid name you entered for node 1.
  - **Grid Master's IP Address:** Enter the same VIP you entered for node 1.
  - **Shared Secret:** Enter a text string that both appliances use as a shared secret to authenticate each other when establishing a VPN tunnel between them. This must match your entry in node 1.
5. On the next screen verify the IP address settings of the member and click **Next**.
6. The last screen displays the settings you specified in the previous panels of the wizard. Verify that the information is correct and click **Finish**.  
To confirm that the appliance has successfully joined the Grid, log in to the Grid Master and navigate to the **Grid** tab, select the **Grid Manager** -> **Members** tab. This panel lists the Grid members. Check the icon in the Status column of the newly added member. (Green = The appliance has joined the Grid and is functioning properly; Yellow = The appliance is in the process of joining the Grid; Red = The appliance has not joined the Grid). You can also use the CLI command set network to join an appliance to a Grid.

## Grouping Members by Extensible Attributes

When you have a few members in your Grid, you can organize and group them by extensible attributes that contain the same values. Using the **Group Results** function, you can organize your members in a meaningful way and quickly identify them based on common data. When you group members by multiple extensible attributes, the appliance groups the members hierarchically based on the order of the filters. For example, when you filter members first by extensible attribute "Site equals London" and then by extensible attribute "Organization equals Engineering," the appliance groups corresponding members first by Site and then by Organization based on the values you enter. In the **Grid** tab -> **Grid Manager** tab -> **Members** tab, Grid Manager displays the grouped members in a hierarchical view that displays the member group name (**London**). You can click the London link and drill down to the next level of grouping. In this case, Grid Manager displays the organization group (**Engineering**) in the **Members** tab. When you click the **Engineering** link to drill down to the next level, all associated members that belong to this member group (**London -> Engineering**) are displayed.

To go back to a previous hierarchical view, click the link of the corresponding level in the breadcrumb.

To group members by extensible attributes:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.  
or  
From the **Grid** tab, select the **Grid Manager** tab -> **Services** tab.  
Note that you can use the "Group Results" function for the following services: DNS, DHCP, TFTP, FTP, HTTP, NTP, bloxTools, Captive Portal, and Reporting services.  
or  
From the **Data Management** tab, select the **DHCP**, **File Distribution**, or **DNS** tab -> **Members/Servers** tab.
2. Complete the following to group members with the same extensible attribute value:
  - **Group Results:** Select this checkbox to enable the appliance to group members by extensible attributes.
  - **Group By:** From the drop-down list, select the first extensible attribute that you want the appliance to use for filtering members.

Grid Manager displays data per group of members configured with the same extensible attribute value.

To add additional **Group By** filter, click the + icon, and then select a value from the drop-down list. You can apply up to 10



**Group By** filters. You can also delete a filter by clicking the - icon.

When you enable reporting service on the Grid and configure multi-site cluster, you can group reporting members by reporting site extensible attributes. For information about reporting clusters, see [Configuring Reporting Clusters](#). Grid Manager displays the following information for the specified extensible attribute:

- *<Selected extensible attribute>*: Displays the selected extensible attribute value.
- **Status**: This is the overall status for all members in the group. Depending on the status of each member, the overall status can be one of the following: **Working**, **Warning**, **Failed**, **Offline**, **Inactive**, or **Unknown**. For information about the status, see [Status Dashboard](#).

Note that in an HA pair, when one of the appliances is in the **Working** status and the other appliance has a status other than **Working**, **Inactive**, and **Unknown**, then the overall status of HA members is **Warning**. When you use filters and the group by extensible attribute feature, filters take precedence over the group by function.

When you drill-down to the member level, Grid Manager displays the members in the group.

## Configuring an IPv6-only Grid

An IPv6-only Grid uses IPv6 as the communication protocol and it includes an IPv6 Grid Master and the Grid members, which can be either IPv6 or dual mode (IPv4 and IPv6). In order to configure an IPv6-only Grid, you have to first create an IPv6 Grid Master and then join the Grid members using their IPv6 addresses.

The process of configuring an IPv6-only Grid involves the following steps:

1. Make a console connection to the Grid Master and configure an IPv6 address for the Grid Master using the CLI command `set network`. For information, see [Method 2 – Using the CLI](#).
2. Open a web browser and make an HTTPS connection to the IPv6 address of LAN1 port of the Grid Master.
3. Log in using the default username and password **admin** and **infoblox**. For detailed information about logging in to the GUI, see [Logging on to the NIOS UI](#).
4. The Grid Setup Wizard appears when you first log in to the appliance. You can use it to create an IPv6 HA Master or IPv6 single Grid Master. For information about creating an HA Master, see [Creating a Grid Master#CreatinganHAGridMasterCreating an HA Grid Master](#) and for information about creating a single Grid Master, see [Creating a Single Grid Master](#).  
The type of network connectivity for the Grid Master should be set to IPv6. To verify, navigate to the **Grid** tab -> **Grid Manager** tab -> **Members** tab -> *member* checkbox -> Edit icon. In the *Grid Member Properties* editor, select the **Network** tab -> **Basic** tab, check that the **Type of Network Connectivity** is set to **IPv6**. The Grid members can join the Grid Master using IPv6 only.  
Note:  
You can add additional IPv4 and IPv6 addresses for LAN2 and MGMT ports for the Grid services, in the **Additional Ports and Addresses** table. But for an IPv6-only Grid, you can configure IPv6 address for the VLAN port.
5. Legacy Data Connector virtual machines are not supported on IPv6-only Grids.
6. Add IPv6 single members and HA members to the Grid. For information, see [Adding Grid Members](#).  
Note that to add a discovery member to an IPv6-only grid, add the member first and then install its discovery license. If the system displays the "Cannot configure IPv6-only settings on member because it has a discovery license installed." in the **Grid Member Properties Editor** dialog box of the discovery member, disregard the error message.
7. You can use the Grid Setup Wizard or access the *Join Grid* dialog box to join appliances to a Grid. See [Joining Appliances to the Grid](#).

You can also configure IPv6 address for the MGMT interface of the appliance and join the Grid through the MGMT interface.

## Transforming to an IPv6-only Grid

After an upgrade from a previous NIOS release to NIOS 7.0 or later, each node in the Grid is set to either IPv4 or dual mode (IPv4 and IPv6). Transforming an IPv4-only Grid to an IPv6-only Grid may take a longer duration. Hence, before removing the IPv4 addresses from each Grid member, you have to configure additional IPv6 Grid communication protocol for each Grid member so that all the services function properly using IPv6.

Note the following before converting an IPv4-only or a dual mode Grid to an IPv6-only Grid:

- If a Grid member is designated as a Master Candidate, the Grid manager does not allow you to change the type of network connectivity of the Grid Master or the Grid Master Candidate. Therefore, you must disable the Grid member from being a Grid Master Candidate before changing the type of network connectivity of the Grid Master or the Grid Master Candidate. You can deselect the **Master Candidate** option in the **General** tab of the *Grid Member Properties* editor to disable a member from being a Master Candidate. Note that at this point, the Grid will not have a Grid Master Candidate and this may result in an unrecoverable condition, if the Grid Master goes down.
- You must stop all services on the Grid Master and Grid members that uses IPv4.
- If external servers like authentication servers, forwarders, root name servers, backup servers, etc. is configured with IPv4 addresses, then it will not work after converting the Grid to an IPv6-only Grid. Hence, make sure that you change the IPv4 address of the external server to IPv6 address before converting the Grid to an IPv6-only Grid.

Note that Infoblox recommends that you back up the configuration after you convert a Grid to IPv6-only mode. Restoring the old backup by performing a forced restore, may prevent the Grid members from rejoining the Grid Master after the restore.

The process of transforming an IPv4-only or dual mode Grid to an IPv6-only Grid involves the following steps:

1. Convert the Grid Master into dual mode (IPv4 and IPv6) if it is in IPv4 mode, as follows:
  - Login to the Grid Master, from the **Grid** tab -> **Grid Manager** tab -> **Members** tab -> select the Grid Master and click the Edit icon.
  - In the *Grid Member Properties* editor, select the **Network** tab -> **Basic** tab.
  - In the **Type of Network Connectivity** field, select **IPv4 and IPv6** from the drop-down list and enter the network information for LAN1 (IPv6) address in the **Ports and Addresses** table.  
For HA Master, select **IPv4** in the **Send HA and Grid Communication Over** field, and enter the network information for VIP (IPv6), Node1 LAN1 (IPv6), Node2 LAN1 (IPv6) in the **Ports and Addresses** table.
  - Save the configuration and click **Restart** if it appears at the top of the screen.
2. Similarly, convert all the Grid members into dual mode (IPv4 and IPv6) if it is in IPv4 mode. All the members will rejoin the Grid using IPv4.
3. Force each Grid member to rejoin the Grid using IPv6, as follows:
  - From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox -> Edit icon.
  - In the *Grid Member Properties* editor, select the **Network** tab -> **Basic** tab.
  - In the **Always Force or Prefer this Communications Protocol** field, select **IPv6** from the drop-down list, and also select **IPv6** in the **Send HA and Grid Communication over** field if the member is an HA pair.  
This setting will force the Grid member to rejoin the Grid using IPv6 and it uses IPv6 for all the services.
  - Save the configuration and click **Restart** if it appears at the top of the screen.
4. Configure each Grid member to provide DNS service using IPv6, as follows:
  - From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox -> Edit icon.
  - In the *Member DNS Properties* editor, select the **General** tab -> **Basic** tab.
  - Enable the IPv6 checkbox for the desired interface (LAN1, LAN2, or MGMT) under **DNS Interfaces**.
  - Ensure that the primary and secondary servers are configured with an IPv6 address for each zone, before disabling the IPv4 checkbox for LAN1, LAN2, or MGMT interface.
  - Save the configuration and click **Restart** if it appears at the top of the screen.

Note:

When you transform an IPv4-only or dual mode Grid to an IPv6-only Grid, the LAN1 port for IPv4 is always enabled. The LAN1 port is disabled only when the Grid is configured using IPv6 from the beginning.

  - For vNIOs appliances, some of the options in the *DNS Interfaces* section may vary depending on your vNIOs configuration. For example, if you are using a single network interface instance of vNIOs for GCP, you will see choices specific to the LAN1 interface only. For more information, see the vNIOs documentation specific to your product at [Appliances](#).
5. Configure each Grid member to provide DHCP service using IPv6, as follows:
  - From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox -> Edit icon.
  - In the *Member DHCP Properties* editor, select the **General** tab -> **Basic** tab.
  - Disable the IPv4 checkbox for LAN1 and LAN2 interface under **DHCP Interfaces**.
  - Enable the IPv6 checkbox for the desired interface (LAN1 or LAN2) under **DHCP Interfaces**.  
If IPv6 network is not configured, you can create an IPv6 network. For information, see [Managing IPv6 Networks](#).
  - Save the configuration and click **Restart** if it appears at the top of the screen.

6. Enable each Grid member and the Grid Master to use IPv6 for all the services, as follows:
  - From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox -> Edit icon.
  - In the *Grid Member Properties* editor, select the **Network** tab -> **Basic** tab, and select **Customized Settings** and specify the following:
    - **Always Force this Communication Protocol for:** Select **IPv6** from the drop-down list for the Grid and reporting service.
    - **Always Prefer this Communication Protocol for:** Select **IPv6** from the drop-down list as the preferred communication protocol for the listed services which has two types of resolution (A and AAAA records). The appliance uses the preferred protocol first for the service.
    - Save the configuration and click **Restart** if it appears at the top of the screen.
7. When all the services are functioning using IPv6, you can remove the IPv4 addresses from all the Grid members by converting the Grid member from dual mode (IPv4 and IPv6) to IPv6 mode, as follows.
  - From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox -> Edit icon.
  - In the *Grid Member Properties* editor, select the **Network** tab -> **Basic** tab.
  - In the **Type of Network Connectivity** field, select **IPv6** from the drop-down list.
  - Save the configuration and click **Restart** if it appears at the top of the screen.
8. Similarly, you can convert the Grid Master from dual mode (IPv4 and IPv6) to IPv6 mode after converting all the Grid members to IPv6 mode.  
Ensure to remove the IPv4 addresses from the Grid Master only after you remove the IPv4 addresses from all the Grid members.

## Auto-Provisioning NIOS Appliances

In addition to using the Grid Setup Wizard or access the *Join Grid* dialog box to join appliances to a Grid, you can set up an appliance using the auto-provisioning feature, which allows a DHCP server to automatically assign an IP address to the appliance. You can then join the auto-provisioned appliance to the Grid.

Auto-provisioning is enabled by default for new installations of physical appliances, but it is not supported for vNIOS appliances. When you connect the appliance to the network, a lease request is automatically sent to the DHCP server. The DHCP server fingerprints the client as "Infoblox Appliance," as the DHCP client provides the unique option sequence (1,28,2,2,3,3,15,6,12) and vendor ID (INFOBLOX). The DHCP server assigns a DHCP lease and a dynamic IP address to the appliance. If the DHCP lease request fails, the default IP address (192.168.1.2) is assigned to the appliance. The DHCP client tries to send the lease request for a duration of one minute when the appliance is either in the factory default state or in the auto-configured default IP address state after a reboot. If you do not use auto-provisioning to set up the appliance, then you can wait one minute before connecting the appliance to the network. Otherwise, the DHCP server will assign a dynamic IP address to the appliance. Note that if you have already set the IP address for the appliance through the Infoblox CLI, GUI, or API, then auto-provisioning is disabled for the appliance and the lease address is not requested. When auto-provisioning is enabled for an appliance, the DNS, DHCP, FTP, TFTP, HTTP, NTP, bloxTools, Captive Portal, Reporting services, as well as backup and restore are disabled for the member until a static IP address is set for the appliance. You can join a single appliance or HA pair to the Grid. After the appliance joins the Grid, the static IP address is set for the appliance.



### Note

- Auto-provisioning supports only IPv4 addressing and not IPv6 addressing.
- Auto-provisioning supports only physical appliances. You cannot auto-provision a hybrid HA setup of a physical and virtual appliance.

When you upgrade or downgrade the appliance to a release that includes this feature, auto-provisioning is disabled for the appliance.

## Setting Up Physical Appliances Using Auto-Provisioning

Complete the following to set up an appliance using auto-provisioning and to join the auto-provisioned appliance to the Grid Master:

1. Connect the appliance to a network using an Ethernet cable, connect it to a power source, and then turn on the power. For information about cabling the appliance to a network and powering the appliance, refer to the user guide or installation guide that ships with the product.  
A lease request is automatically sent to the DHCP server that assigns a DHCP lease and a dynamic IP address to the appliance. The DHCP client tries to send the lease request for a duration of one minute and if the request fails, the default IP address (192.168.1.2) is assigned to the appliance.
2. Join the appliance to the Grid Master. You can join the auto-provisioned appliance to the Grid Master using the *Connect* dialog box. You can also join an appliance to the Grid using the *JoinGrid* dialog box. For more information, see [Adding Grid Members](#) *Joining Appliances to the Grid*.  
A static IP address is set and auto-provisioning is automatically disabled for the appliance after it joins the Grid.

Only the following physical platforms supports Auto-Provisioning:

IB-4015, IB-4025, TE-815, TE-825, TE-1415, TE-1425, TE-2215, TE-2225, TR-1405, TR-2205, TR-4005.



#### Note

When auto-provisioning is disabled for an appliance and the network address is not preserved, auto-provisioning will be re-enabled and a DHCP lease request is sent to the DHCP server if you reset the appliance using the CLI command `reset all auto_provision` or reset the database using the CLI command `reset database auto_provision`. However, if the static IP address for an appliance is set and network settings are preserved, auto-provisioning will be re-enabled for the appliance but the lease address will not be requested if you reset the database using the CLI command `reset database auto_provision`.

## Joining Auto-Provisioned Appliances to the Grid

You can join a predefined appliance with a DHCP lease to the Grid Master using the *Connect* dialog box. You can join a single appliance or an HA pair to the Grid Master. For an HA pair, the member which is offline will join the Grid Master and it will become the active node. When both the members of an HA pair are offline, Node 1 of an HA pair is joined to the Grid Master.

Only superusers can join a Grid member to the Grid Master. If the Grid member fails to join the Grid, then the remote console is enabled for the appliance and you can join the appliance to a Grid through the remote console. You can log in to the remote console using the user name, **admin** and the Grid shared secret as the password.

To join a single appliance or an HA pair to a Grid Master, complete the following:

1. Log in to the Grid Master. Note that the single appliance or the HA pair must be online and the Grid Master must be able to reach the appliance.
2. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
3. Add the appliance as a Grid member. For information about adding Grid members to the Grid, see [Adding Grid Members](#).
4. Select the Grid member that you want to join to the Grid Master, expand the Toolbar and click **Connect**.
5. The following fields are displayed in the *Connect* dialog box:
  - **Host Name:** The name of the member.
  - **Configured IPv4 Address:** The IPv4 address of the member.
  - **Site:** The site to which the IP address belongs. This is one of the predefined extensible attributes.
  - **Temporary IPv4 Address:** Enter the IPv4 address of the DHCP lease or click **Select** to select the DHCP lease.  
Grid Manager displays the *Lease Selector* dialog box from which you can select the DHCP lease. Note that the *Lease Selector* displays the active DHCP leases which are fingerprinted as "Infoblox Appliance".
6. Click **Next** to retrieve the appliance information.  
The Grid Master uses SSL to connect to the appliance and it gets the appliance information. Grid Manager displays the following information for the appliance:
  - **Remote Appliance Type:** The appliance type.
  - **Remote Appliance Serial Number:** The serial number of the appliance.

- **Licenses:** Grid Manager displays the Grid license and the licenses that are pre-provisioned on the member.  
It displays the following information:
    - **Type:** The license type.
    - **String:** The license string. If the license string is not displayed, you can enter or paste it here.
7. Click **Connect** to join the appliance to the Grid Master.  
To confirm that the appliance has successfully joined the Grid, check the status in the **Status** column of the newly added member. (Green = The appliance has joined the Grid Master and is functioning properly; Yellow = The appliance is in the process of joining the Grid Master; Red = The appliance has not joined the Grid Master).

## Pre-Provisioning NIOS and vNIOS Appliances

Before joining a member to the Grid, you can first enable provisional licenses and make necessary configurations on the offline member, which allows DNS and DHCP data to be associated with the member prior to its deployment. Note that pre-provisioned members are treated as offline members. There are a few guidelines to consider before you pre-provision a member. For more information about the guidelines, see [Guidelines for Pre-provisioning Offline Grid](#). When you add a new member to the Grid, the **Pre-Provisioning** tab is displayed in the *Grid Member Properties* editor. You can pre-provision the member by defining its hardware model and enable certain provisional licenses through the **Pre-Provisioning** tab. This tab is not displayed after the member successfully joins the Grid. NIOS supports the following provisional licenses: Cloud Platform, DHCP, DNS, DNS Traffic Control, Enterprise (formerly Grid), FireEye, Microsoft Management, RPZ (Response Policy Zone), and vNIOS. You must enable provisional licenses before you can make supported configurations on the pre-provisioned member. For more information about these licenses, see [About Provisional Licenses](#).



### Note

You must have the Enterprise and vNIOS licenses pre-provisioned in order for a vNIOS appliance to join the Grid. For a cloud virtual appliance, include the Cloud Platform license.

To pre-provision an offline Grid member and join it to the Grid at a later time, complete the following:

1. Add a new single member or HA member to the Grid, as described in [Adding a Single Member](#) or [Adding an HA Member](#).
2. Pre-provision the offline member, as described in [Configuring Pre-Provisioned Members](#) below.
3. Configure services to use the pre-provisioned member.
4. Obtain permanent licenses you have specified for pre-provisioning and use the **set license** CLI command to install the licenses on the member. For more information about CLI commands, refer to the *Infoblox CLI Guide*.
5. Join the pre-provisioned member to the Grid, as described in [Joining Appliances to the Grid](#). For guidelines about joining pre-provisioned members, see [Guidelines for joining Pre-Provisioned Members](#).

## Guidelines for Pre-provisioning Offline Grid Members

Before you pre-provision a Grid member, consider the following:

- A pre-provisioned Grid member is an offline member. When you upgrade a Grid that has a pre-provisioned member, the upgrade behaves the same way as it does when you upgrade the Grid that has an offline member. Note that you cannot pre-provision a member or update its settings during a scheduled upgrade. For more information about upgrades, see [Upgrading NIOS](#).
- You cannot change the pre-provisioned member configuration after you save it. To change the configuration, you must first delete the member and pre-provision it again. If you want to delete certain provisional licenses or change the hardware model for the pre-provisioned member, you must also first delete the existing member and define a new one. For information about deleting a member, see [Removing a Grid Member](#).
- When you assign a network, zone, or IPv4 DHCP failover association to a pre-provisioned member, the **Restart Service** button is not displayed. If you restart any service on a pre-provisioned member, no action is actually taken even though you may receive a message indicating that the operation may take a few minutes. When you



join the member to the Grid, NIOS will run respective member services on the joined member. For more information about service restarts, see [Restarting Services](#).

- NIOS allows you to backup information about the pre-provisioned member. When you perform a forced restore however, NIOS does not restore the pre-provisioned licenses if you have already installed permanent NIOS licenses on the corresponding member. For more information about backup and restore, see [Backing up and Restoring Configuration Files](#).
- You can use **Manage Member Services** to manage the pre-provisioned member services. For more information, see [Monitoring Member Services](#).

## Configuring Pre-Provisioned Members

The pre-provisioning feature is disabled by default. You must select a supported hardware model for the member to enable this feature.

To pre-provision an offline member, login to the Grid Master and complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click the Add icon.
2. In the *Add Grid Member* wizard, add a new member as described in [Adding a Single Member](#).
3. After you add the member to the Grid, select the member in the **Members** tab and click the Edit icon.
4. In the *Grid Member Properties* editor, select the **Pre-Provisioning** tab, and complete the following:
  - **Member Type**: Displays the member type that you have selected in the **General** tab. The pre-provisioning feature is supported only for **Infoblox** and **Virtual NIOS** member types. Note that you must select a hardware model for the member in order to enable the pre-provisioning feature.
  - **Hardware Model**: Select the hardware model from the drop-down list. Grid Manager displays only the supported hardware models for the specified member type. Once you select the hardware model, the pre-provisioning feature is enabled for the member. NIOS allows you to pre-provision HA members that have the same or different hardware models for Node 1 and Node 2. A few hardware specific features, such as DSCP, VLAN, LAN2, and LOM (Light Out Management), are enabled based on the pre-provisioned hardware model you specify here.
  - **Provisional Licenses**: Select the licenses that you want to enable for the pre-provisioned member. You can select the licenses only after you have specified the hardware model for the member. Once you select and enable a license, you can no longer modify the hardware model for the member. Note that the permanent licenses that you later add to the member must include the ones that are specified for pre-provisioning.
5. Save the configuration.

Note that after you save the configuration, you can no longer modify the hardware model for the member. You also cannot disable any provisional licenses, though you can add new ones. To disable provisional licenses, you must first remove the pre-provisioned member and then configure a new one.

## About Provisional Licenses

If a member has never joined a Grid, you can pre-provision this member provided that you define the hardware model for the member and assign provisional licenses to it. Provisional licenses are not permanent NIOS licenses. Though they do not have expiration dates or validity periods, you must replace these licenses with corresponding permanent licenses before you join the member to the Grid.



### Note


Before you join the member to the Grid, use the CLI command `set license` to add corresponding permanent licenses that you have specified for pre-provisioning. For information about CLI commands, refer to the Infoblox CLI Guide{ }. You can also allocate pre-purchased licenses from the pool. For information, see You can use the following OpenStack cloud-init template to configure an IB-V815 as a Grid Master

NIOS supports the following provisional licenses: Cloud Platform, DHCP, DNS, DNS Traffic Control, Enterprise (formerly Grid), FireEye, Microsoft Management, RPZ (Response Policy Zone), and vNIOS.

After you configure the offline member, you can select the pre-provisioned member from the corresponding wizards and editors based on the required license(s). The following table lists the wizards and editors from which you can select a pre-provisioned member when required pre-provisioned licenses are enabled:

Wizards and editors from which you can select a pre-provisioned member	Required license(s)
DNS Zones and Name Server Groups	dns
DHCP IPv4 and IPv6 networks	dhcp
IPv4 DHCP Failover Association	dhcp
Microsoft servers Note that the initial synchronization with Microsoft servers is read-only. When you join the appliance to the Grid, the appliance removes all Microsoft management objects that you have configured on the Microsoft servers after the synchronization. The configuration on the Microsoft servers will replace the configuration on the NIOS appliance.	ms_management
Grid Members	vnios
Grid license	enterprise
Response Policy Zones	RPZ
Response Policy Zones	FireEye
DNS Traffic Control	DTC
Cloud tabs and related wizards and editors	cloud_api

## Note

 If you configure a DHCP Failover using an online member and a pre-provisioned member, assign it to a range, and start DHCP service, no addresses will be served because the initial synchronization does not happen due to the pre-provisioned offline member. NIOS logs the following message in the syslog:

```
2013-12-24T08:37:23+00:00 daemon (none) dhcpd[8790]: info DHCPDISCOVER
from cb:86:a8:45:6c:5c via 10.120.21.236: not responding (recovering)
```

## Guidelines for Joining Pre-Provisioned Members to the Grid

Before you join a pre-provisioned member to the Grid, ensure that you verify the appliance model and provisional licenses for the member. For information about how to join a member to the Grid, see [Joining Appliances to the Grid](#). Note the following about joining a pre-provisioned member to the Grid:

- If you install fewer permanent licenses than the specified provisional licenses, you cannot join the member to the Grid.
- If the pre-provisioned member does not have any provisional licenses enabled, you can join the member to the Grid provided that you install a permanent Grid license on the member.
- You must install at least the set of permanent licenses that were specified for pre-provisioning along with any other needed licenses, except for the following:

- You can join the member to the Grid if the pre-provisioned member is a vNIOS virtual appliance and has only the DNS license enabled, and you install both the vNIOS and DNS licenses on the member.
- Similarly, you can join the member to the Grid if the pre-provisioned member is a vNIOS virtual appliance and has both DNS and DHCP licenses enabled, and you install the vNIOS, DNS, and DHCP licenses on the member.
- After you successfully join the pre-provisioned member to the Grid, provisional licenses are removed and permanent licenses take effect.
- After the member joins the Grid successfully, the **Pre-Provisioning** tab is not displayed in the *Grid Member Properties* editor.

## Configuring a Grid

You can configure an IPv4-only, IPv6-only, or a dual mode (IPv4 and IPv6) Grid, but the configuration example uses IPv4 addresses. In this example, you configure seven NIOS appliances in a Grid serving internal DHCP and DNS for an enterprise with the domain name corpxyz.com. There are four sites: HQ and three branch offices. A hub-and-spoke VPN tunnel system connects the sites, with HQ at the hub. The distribution and roles of the NIOS appliances at the four sites are as follows:

- HQ site (four appliances in two HA pairs):
  - HA Grid Master: Hidden primary DNS server
  - HA member: Secondary DNS server and DHCP server for HQ
- Site 1 (two appliances in an HA pair): HA member, secondary DNS server and DHCP server for Site 1
- Site 2 (one appliance): Single member, secondary DNS server and DHCP server for Site 2



### Note

When adding an Infoblox appliance to an existing Grid, you must first check whether the Grid is running the minimum required software release of the appliance. For information, refer to the document, *Minimum Required Release Software for Hardware Platforms*, that was shipped with your product.

To create a Grid, you first create a Grid Master and then add members. The process involves these three steps:

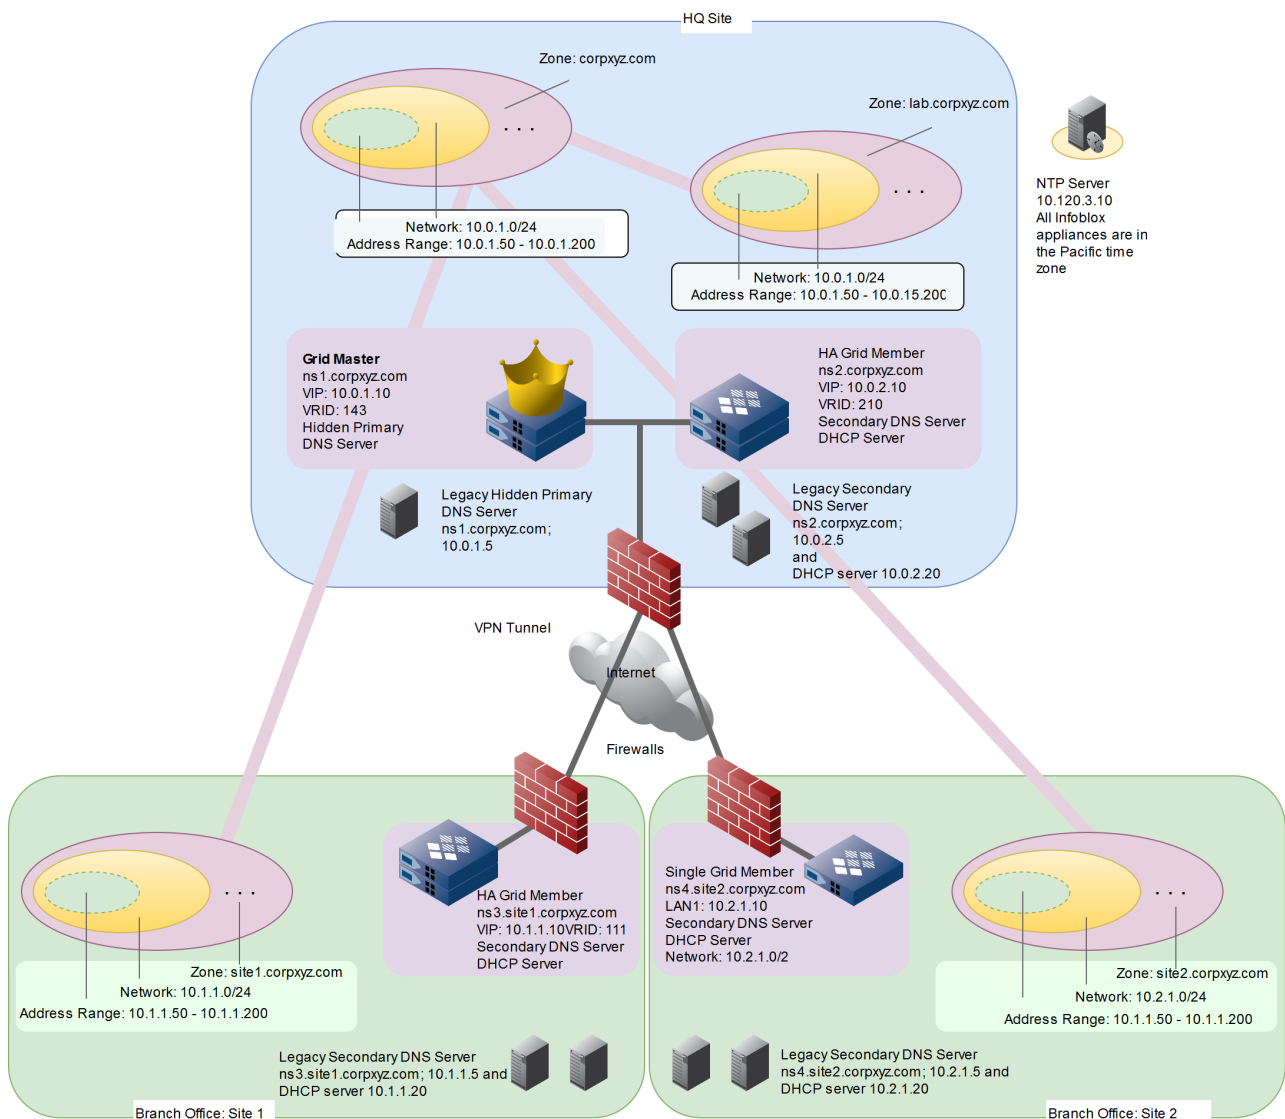
1. Configuring two appliances at HQ as the Grid Master. For more details, see [Create the Grid Master](#) below.
2. Logging in to the Grid Master and defining the members that you want to add to the Grid; that is, you configure Grid member settings on the Grid Master in anticipation of later joining those appliances to the Grid. For more details, see [Define Members on the Grid Master](#) below.
3. Logging in to the individual appliances and configuring them so that they can reach the Grid Master over the network and join the Grid. For more details, see [Join Appliances to the Grid](#) below.

After creating the Grid and adding members, you use the Data Import Wizard to import DHCP and DNS data from legacy servers. For more details, see [Import DHCP Data](#) and [Import DNS Data](#) below.

Finally, you transition DHCP and DNS service from the legacy servers to the Infoblox Grid members. For more details, see [Enable DHCP and Switch Service to the Grid](#) below.

*Network Diagram*





## Cable All Appliances to the Network and Turn On Power

Cable the NIOS appliances to network switches. After cabling each appliance to a switch and connecting it to a power source, turn on the power. For information about installing and cabling the appliance, refer to the user guide or installation guide that ships with the product.

1. At HQ and Site 1, connect Ethernet cables from the LAN1 and HA ports on the appliances in each HA pair to a switch, connect the appliances to power sources, and turn on the power for each appliance.  
Note that when connecting the nodes of an HA pair to a power source, connect each node to a different power source if possible. If one power source fails, the other might still be operative.
2. At Site 2, connect an Ethernet cable from the LAN1 port on the single appliance to a switch, connect the appliance to a power source, and turn on the power for that appliance.

## Creating the Grid Master



### Note

IPv6 addressing is fully supported on Infoblox Grid Masters, HA pairs and standalone HA pairs, and appliances. Examples in the sections of this chapter use IPv4.

Configure two appliances at HQ to be the two nodes that make up the HA pair forming the Grid Master.

### Grid Master – Node 1

1. By using the LCD or by making a console connection to the appliance that you want to make Node 1 of the HA pair for the Grid Master, change the default network settings of its LAN1 port to the following:
  - **IP Address:** 10.0.1.6
  - **Netmask:** 255.255.255.0
  - **Gateway:** 10.0.1.1
2. Connect your management system to the HQ network, open a browser window, and connect to <https://10.0.1.6>.
3. Log in using the default username and password **admin** and **infoblox**.
4. Review the End-User License Agreement and click **I Accept**.
5. Read about the *Infoblox Customer Experience Improvement Program* and choose whether to participate (opt in) or not participate (opt out) in the program. By default, participation is enabled. If you want to opt out of the program, select **To Opt-Out of the alert program, please click here**. For more information about the program, see [Configuring the Customer Experience Improvement Program](#).
6. Click **OK**. The *Grid Setup* wizard appears.
7. On the first screen, select **Configure a Grid Master** and click **Next**.
8. Specify the Grid properties:
  - **Grid Name:** Enter **corpxyz**.
  - **Shared Secret:** Enter **Mg1kW17d**.
  - **Confirm Shared Secret:** Enter **Mg1kW17d**.
  - **Hostname:** Enter **ns1.corpxyz.com**.
  - **Type of Network Connectivity:** Select **IPv4** from the drop-down list.
  - **Is the Grid Master an HA pair?:** Select **Yes**.
9. Specify the network properties and click **Next**:
  - **Virtual Router ID:** Enter **143**.
  - **Required Ports and Addresses:** Enter the details in the table to set up the HA pair, see at the end of the procedure.
10. **Enter a new password:** **1n85w2IF**. Retype it and click **Next**.
11. Then, complete the following:
  - **Time zone:** Select **(UTC – 8:00 Pacific Time (US and Canada), Tijuana)**
  - Enable **NTP**, click the Add icon and enter the IP address of the NTP server: **10.120.3.10**
12. Click **Finish**. When you click **Finish**, the Infoblox GUI application restarts.

Interface	Address	Subnet Mask (IPv4) or Prefix Length (IPv6)	Gateway	Port Setting
VIP ( IPv4)	10.0.1.10	255.255.255.0	10.0.1.1	Automatic
Node2 HA (IPv4)	10.0.1.9	255.255.255.0	10.0.1.1	Automatic
Node2 LAN1 (IPv4)	10.0.1.8	255.255.255.0	10.0.1.1	Automatic
Node1 HA (IPv4)	10.0.1.7	255.255.255.0	10.0.1.1	Automatic

Interface	Address	Subnet Mask (IPv4) or Prefix Length (IPv6)	Gateway	Port Setting
Node1 HA (IPv4)	10.0.1.7	255.255.255.0	10.0.1.1	Automatic

## Grid Master – Node 2

- By using the LCD or by making a console connection to the appliance that you want to make Node 2 of the HA pair for the Grid Master, change the default network settings of its LAN1 port to the following:
  - IP Address:** 10.0.1.8
  - Netmask:** 255.255.255.0
  - Gateway:** 10.0.1.1
- In the login window, type **10.0.1.8** in the **Hostname** field.
- Log in using the default username and password, **admin** and **infoblox**.
- From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox.
- Expand the Toolbar and click **Join Grid** and specify the following:
  - Virtual IP of Grid Master:** 10.0.1.10
  - Grid Name:** Enter **corpxyz**
  - Grid Shared Secret:** Enter **Mg1kW17d**
- Confirm the configuration, and then on the last screen of the wizard, click **Finish**. The HTTPS session terminates, but the login window remains open.
- In the login window, type **10.0.1.10** (the VIP address for the Grid Master) in the **Hostname** field.
- Log in using the default username **admin** and the password **1n85w2IF**.
- To check the status of the two nodes of the HA Grid Master, navigate to the **Grid** tab, select the **Grid Manager** -> **Members** tab. This panel lists the Grid members. Check the icon in the **Status** column of the Grid Master. The colors indicate the following:
  - Green:** The appliance has joined the Grid and is functioning properly.
  - Yellow:** The appliance is in the process of joining the Grid
  - Red:** The appliance has not joined the Grid.

You can also use the CLI command `set network` to join an appliance to a Grid. Check that the status indicators are all green in the *Detailed Status* panel.

During the joining process, an appliance passes through the following four phases:

- Offline:** The state when a Grid member—in this case, the second node of the HA pair composing the Grid Master—is not in contact with the active node of the master.
- Connecting:** The state when an appliance matching a member configuration contacts the master to join the Grid and negotiates secure communications and Grid membership.
- Synchronizing:** The master transmits its entire database to the member.
- Running:** The state when a member is in contact with the master and is functioning properly.



### Note

Depending on the network connection speed and the amount of data that the master needs to synchronize with the member, the process can take from several seconds to several minutes to complete.

## Defining Members on the Grid Master

Before logging in to and configuring the individual appliances that you want to add to the Grid, define them first on the Grid Master. You can configure a Grid member in IPv4, IPv6, or a dual mode (IPv4 and IPv6), but the configuration example uses IPv4 addresses.

### HQ Site – HA Member

- From the **Grid** tab, select the **Grid Manager** -> **Members** tab.
- Expand the Toolbar and click **Add** -> **Add Grid Member**.
- In the *Add Grid Member* wizard, complete the following and click **Next**:

- **Member Type:** Select **Infoblox**.
  - **Host Name:** Enter **ns2.corpxyz.com**.
  - **Comment:** Enter **HQ Site - ns2.corpxyz.com**.
4. Enter the following information about the member that you are adding to the Grid and click **Save & Close**:
- **Type of Network Connectivity:** Select **IPv4** from the drop-down list.
  - **High Availability Pair:** Select this option.
  - **Virtual Router ID:** **210**
  - **Required Ports and Addresses:**

Interface	Address	Subnet Mask (IPv4) or Prefix Length (IPv6)	Gateway	Port Settings
VIP (IPv4)	10.0.2.10	255.255.255.0	10.0.2.1	Automatic
Node1 HA (IPv4)	10.0.2.7	255.255.255.0	10.0.2.1	Automatic
Node2 HA (IPv4)	10.0.2.9	255.255.255.0	10.0.2.1	Automatic
Node1 LAN1 (IPv4)	10.0.2.6	255.255.255.0	10.0.2.1	Automatic
Node2 LAN1 (IPv4)	10.0.2.8	255.255.255.0	10.0.2.1	Automatic

#### Site 1 – HA Member

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Expand the Toolbar and click **Add** -> **Add Grid Member**.
3. In the *Add Grid Member* wizard, enter the following and click **Next**:
  - **Member Type:** Select **Infoblox**.
  - **Host Name:** Enter **ns3.site1.corpxyz.com**
  - **Comment:** Enter **Site 1 - ns3.site1.corpxyz.com**
4. Specify the following information about the member that you are adding to the Grid and click **Save & Close**:
  - **Type of Network Connectivity:** Select **IPv4** from the drop-down list.
  - **High Availability Pair:** Select this option.
  - **Virtual Router ID:** Enter **111**.
  - **Required Ports and Addresses:**

Interface	Address	Subnet Mask (IPv4) or Prefix Length (IPv6)	Gateway	Port Settings
VIP (IPv4)	10.1.1.10	255.255.255.0	10.1.1.1	Automatic
Node1 HA (IPv4)	10.1.1.7	255.255.255.0	10.1.1.1	Automatic
Node2 HA (IPv4)	10.1.1.9	255.255.255.0	10.1.1.1	Automatic
Node1 LAN1 (IPv4)	10.1.1.6	255.255.255.0	10.1.1.1	Automatic
Node2 LAN1 (IPv4)	10.1.1.8	255.255.255.0	10.1.1.1	Automatic

#### Site 2 – Single Member

1. From the **Grid** tab, select the **Grid Manager** -> **Members** tab.
2. Expand the Toolbar and click **Add** -> **Add Grid Member**.
3. In the *Add Grid Member* wizard, enter the following and click **Next**:

- **Member Type:** Select **Infoblox**
  - **Host Name:** **ns4.site2.corpxyz.com**
  - **Comment:** **Site 2- ns4.site2.corpxyz.com**
4. Specify the following information about the member that you are adding to the Grid and click **Next**:
    - **Type of Network Connectivity:** Select **IPv4** from the drop-down list.
    - **Standalone Member:** Select this option.
    - **Required Ports and Addresses:** Click the empty fields and enter the following information:
      - **Address:** Enter **10.2.1.10**.
      - **Subnet Mask (IPv4) or Prefix Length (IPv6):** Enter **255.255.255.0**
      - **Gateway:** Enter **10.2.1.1**
      - **Port Settings:** Select **AUTOMATIC**.
  5. Save the configuration and click **Restart** if it appears at the top of the screen.
  6. Log out from the Grid Master.

### Join Appliances to the Grid

To complete the process of adding appliances to the Grid, log in to and configure each individual appliance so that it can contact the Grid Master.

#### HQ Site – HA Grid Member (Node 1)

Make a console connection to the appliance that you want to make Node 1 in the HA pair, and then enter the following:

```
Infoblox > set network
```

NOTICE: All HA configuration is performed from the GUI. This interface is used only to configure a standalone node or to join a Grid.

```
Enter IP address: 10.0.2.6
```

```
Enter netmask : 255.255.255.0
```

```
Enter gateway address : 10.0.2.1
```

```
Configure IPv6 network settings? (y or n):n Become Grid member? (y or n): y
```

```
Enter Grid Master VIP: 10.0.1.10
```

```
Enter Grid Name: corpxyz
```

```
Enter Grid Shared Secret: Mg1kW17d
```

```
New Network Settings: IP address: 10.0.2.6
```

```
Netmask: 255.255.255.0
```

```
Gateway address: 10.0.2.1
```

```
Join Grid as member with attributes: Grid Master VIP: 10.0.1.10
```

```
Grid Name: corpxyz
```

```
Grid Shared Secret: Mg1kW17d
```

```
WARNING: Joining a Grid will replace all the data on this node!
```

```
Is this correct? (y or n): y
```

```
Are you sure? (y or n): y
```

The Infoblox application restarts. After restarting, the appliance contacts the Grid Master and joins the Grid as Node 1.

#### HQ Site – HA Member (Node 2)

Make a console connection to the appliance that you want to make Node 2 in the HA pair, and then enter exactly the same data you entered for Node 1 except that the IP address is 10.0.2.8.

After the application restarts, the appliance contacts the Grid Master and joins the Grid as Node 2, completing the HA member configuration for the HQ site.

#### Site 1 – HA Grid Member (Node 1)

Make a console connection to the appliance that you want to make Node 1 in the HA pair at Site 1 and use the `set network` command to configure its basic network and Grid settings. Use the following data:

- **IP Address:** 10.1.1.6
- **Netmask:** 255.255.255.0
- **Gateway:** 10.1.1.1
- **Grid Master VIP:** 10.0.1.10
- **Grid Name:** corpxyz
- **Grid shared secret:** Mg1kW17d

The Infoblox application restarts. After restarting, the appliance contacts the Grid Master and joins the Grid as Node 1.

#### Site 1 – HA Grid Member (Node 2)

Make a console connection to the appliance that you want to make Node 2 in the HA pair at Site 1 and enter exactly the same data that you entered for Node 1 except that the IP address is 10.1.1.8.

After the application restarts, the appliance contacts the Grid Master and joins the Grid as Node 2, completing the HA member configuration for Site 1.

#### Site 2– Single Grid Member

Make a console connection to the appliance that you want to make Node 1 in the HA pair at Site 1 and use the `set network` command to configure its basic network and Grid settings. Use the following data:

- **IP Address:** 10.2.1.10
- **Netmask:** 255.255.255.0
- **Gateway:** 10.2.1.1
- **Grid Master VIP:** 10.0.1.10
- **Grid name:** corpxyz
- **Grid shared secret:** Mg1kW17d

The Infoblox application restarts. After restarting, the appliance contacts the Grid Master and joins the Grid.

To check the status of all the Grid members, log in to the Grid Master at 10.0.1.10, and from the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, select **10.0.1.10** and click the Detailed Status icon. Check that the status indicators are all green in the Detailed Status panel. As an appliance joins a Grid, it passes through the following phases: Offline, Connecting (Downloading Release from Master), Synchronizing, and Running.



#### Note

Depending on the network connection speed and the amount of data that the master needs to synchronize with the member, the process of joining a Grid can take from several seconds to several minutes to complete.

The Grid setup is complete.

## Import DHCP Data

The Data Import Wizard is a software tool that you can download from the Infoblox Support site to your management system. With it, you can import data from legacy DHCP and DNS servers to NIOS appliances. In this example, you use it to import both DHCP and DNS data to the Grid Master at 10.0.1.10, which then uses the database replication mechanism to send the imported data to other Grid members. In the wizard, you also specify which Grid members serve the imported data. The wizard supports various types of DHCP formats, such as the following:

- ISC DHCP
- Lucent VitalQIP
- Microsoft
- Nortel NetID
- CSV (comma-separated values); you can also import IPAM data in CSV format. In this example, all the DHCP data is in standard ISC DHCP format.

## Importing DHCP Data for HQ and Site 2

1. Save the DHCP configuration file from your legacy DHCP server at 10.0.2.20 to a local directory.
2. Visit <https://support.infoblox.com>, log in with your support account, and download the *Data Import Wizard*. The *Data Import Wizard* application downloads to a container within a Java sandbox on your management system and immediately launches, displaying the *Welcome* page.
3. After reading the information in the left panel, click **Next**.
4. Select Import to Infoblox Appliance, enter the following, and then click **Next**:
  - **Hostname or IP address: 10.0.1.10**
  - **Username: admin**
  - **Password: 1n85w2IF**
5. Select the following, and then click **Next**:
  - What kind of data would you like to import? **DHCP/IPAM**
  - Which legacy system are you importing from? **ISC DHCP**
  - Which appliance will be serving this data? **10.0.2.10**
6. Type the path and file name of the DHCP configuration file saved from the legacy server, and then click **Next**.  
Or  
Click **Browse**, navigate to the file, select it, click **Open**, and then click **Next**.
7. In the *Global DHCP Configuration* table, double-click the *Value* cell for the *domain-name-servers* row, and change the IP addresses to **10.0.2.10**.
8. When satisfied with the data, click **Import**.  
You can view the status of the importation process and a summary report in the *Data Import Wizard* Log.
9. To enable DDNS updates, log in to the Grid Master, from the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Properties**.
10. In the **DDNS** -> **Basic** tab of the *Grid DHCP Properties* editor, select **Enable DDNS Updates**.
11. Save the configuration and click **Restart** if it appears at the top of the screen.
12. To check the imported DHCP configuration file, from the **Data Management** tab, select the **DHCP** tab, -> **Members** tab -> **10.0.2.10** checkbox. Expand the Toolbar and click **View DHCP Configuration**.
13. In the DHCP configuration file, check that all the imported subnets are present, and navigate to the beginning of the file and check that you see the

```
ddns-updates on
```

```
statement. (If you see
```

```
ddns-updates
```

```
off, enable DDNS updates for the Grid as explained in steps 9-12.)
```

### Importing DHCP Data for Site 1

1. Repeat the steps as described in Configuring a Grid-Importing DHCP Data for HQ and Site 2, saving the DHCP configuration file from your legacy DHCP server at 10.1.1.20, and importing it to the Grid Master at 10.0.1.10 for the member with IP address 10.1.1.10 to serve.
2. Check the imported DHCP configuration file by logging in to the Grid Master and from the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **10.1.1.10** checkbox. Expand the Toolbar and click **View DHCP Configuration**.

### Importing DHCP Data for Site 3

1. Repeat the steps as described in Configuring a Grid-Importing DHCP Data for HQ and Site 2, saving the DHCP configuration file from your legacy DHCP server at 10.1.1.20, and importing it to the Grid Master at 10.0.1.10 for the member with IP address 10.3.1.10 to serve.
2. After the importation process completes, check the imported DHCP configuration file by logging in to the Grid Master and from the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **10.3.1.10** checkbox. Expand the Toolbar and click **View DHCP Configuration**.

### Import DNS Data

Using the *Infoblox Data Import Wizard*, import DNS data from the legacy hidden primary server at 10.0.1.5 to the new hidden primary server at 10.0.1.10 (the Grid Master). The following are three phases to this task:

1. Before Using the Wizard:
  - Save the named.conf file from the legacy server to a file in a local directory on your management system.
  - Enable the legacy server to perform zone transfers to the NIOS appliance.
  - Configure three name server groups for the Grid and allow the Grid Master/hidden primary DNS server at 10.0.1.10 to receive DDNS updates from the Grid members at 10.0.2.10, 10.1.1.10, and 10.3.1.10. These members act as secondary DNS servers and DHCP servers.
2. Using the Wizard: Define the source, destination, and type of DNS data in the DNS configuration file (named.conf) that you want to import.
3. After Using the Wizard: Check the imported DNS configuration file.

In this example, all the DNS data is in BIND 9 format. The *Data Import Wizard* supports various types of DNS formats, such as the following:

- BIND 4, 8, and 9
- Microsoft
- Lucent VitalQIP
- Nortel NetID

### Before Using the Wizard

You must set up the legacy server and Grid Master before using the *Data Import Wizard*.

### Legacy Server

1. Log in to the legacy name server at 10.0.1.5 and save the named.conf file, which contains all the DNS settings that you want to import into the Infoblox name server, to a local directory on your management system.
2. On the legacy server, enable zone transfers to the NIOS appliance.

### Infoblox Grid Master – DDNS Updates

1. Log in to the Grid Master at 10.0.1.10, and from the **Data Management** tab, select the **DNS** tab -> **Members** tab -> **10.0.1.10** checkbox and select the Edit icon.
2. In the *Member DNS Configuration* editor, select the **Updates** tab -> **Basic** tab and enter the following:
  - Select **Override**.



- Allow updates from: Click the Add icon and select **IPv4 Address**. Enter **10.0.2.10** in the **Name** field of the new row.
3. Click the Add icon again and add **10.1.1.10** and **10.2.1.10** as IP addresses from which you allow DDNS updates.
  4. Save the configuration and click **Restart** if it appears at the top of the screen.  
Note when all DNS servers are members in the same Grid, the members use database replication to synchronize all their data—including DNS zone data. You can change the default behavior so that Grid members use zone transfers instead. In this example, Grid members use database replication.

## Infoblox Grid Master – Name Server Groups

1. From the **Data Management** tab, select the **DNS** tab -> **Name Server Groups** tab.
2. Click the Add icon to open the *Add Name Server Group* wizard.
3. Enter the following:
  - **Name Server Group Name:** HQ-Group
4. Click the Add icon and add the following:
  - **Grid Primary:** ns1.corpxyz.com; **Stealth:** Select this checkbox.
  - **Grid Secondary:** ns2.corpxyz.com; **Grid replication (recommended):** Select this checkbox.
5. Click **Save & New**.
6. Repeat steps 2 to 4 to create another group. Name it **Site1-Group**, and use **ns1.corpxyz.com** as the hidden primary server, **ns3.site1.corpxyz.com** as a secondary server, and Grid replication for zone updates.
7. Repeat steps 2 to 4 to create another group. Name it **Site2-Group**, and use **ns1.corpxyz.com** as the hidden primary server, **ns4.site2.corpxyz.com** as a secondary server, and Grid replication for zone updates.

## Using the Wizard

While progressing through the *Data Import Wizard*, you must define the source, destination, and type of DNS data that you want to import. You then make some simple modifications to the data and import it.

## Defining the Source, Destination, and Type of DNS Data

1. Launch the *Data Import Wizard*.
2. After reading the information in the left panel of the welcome page, click **Next**.
3. Select **Import to Infoblox Appliance**, enter the following, and then click **Next**:
  - **Hostname or IP address:** 10.0.1.10
  - **Username:** admin
  - **Password:** 1n85w2IF

The *Data Import Wizard* Log opens in a separate window behind the wizard. Leave it open while you continue.
4. Select the following, and then click **Next**:
  - **What kind of data would you like to import?:** DNS
  - **Which legacy system are you importing from?:** BIND 9
  - **Which appliance will be serving this data?:** 10.0.1.10
5. Select the following, and then click **Next**:
  - **What BIND 9 DNS configuration file would you like to use?:** Click **Browse**, navigate to the named.conf file you saved from the legacy server, select it, and then click **Open**.
  - **What type of BIND 9 DNS data do you want to import?:** DNS zone information and DNS record data
  - **Where is the BIND 9 DNS record data?:** Zone transfer(s) from a DNS server; 10.0.1.5

The wizard displays two tables of data. The upper table contains global DNS server configuration parameters. The lower table contains zone configurations.

The *Data Import Wizard* Log presents a summary listing the number of views, zones, and DNS records in the configuration file.



### Note

Only superusers can import A, AAAA, shared A, and shared AAAA records with a blank name. Limited-access users must have read/write permission to **Adding a blank A/AAAA record** to import A, AAAA, shared A, and shared AAAA records with a blank name, otherwise, the import operation might fail. You can assign global

permission for specific admin groups and roles to allow to import A, AAAA, shared A, and shared AAAA records with a blank name. For more information, see [Administrative Permissions for Adding Blank A or AAAA Records](#).

## Modifying DNS Data

While importing data from the legacy DNS server, you cancel the importation of global configuration settings, and apply the name server groups you created in [Before Using the Wizard](#) to the zones you want to import.

1. In the *Global DNS Configuration* table, select all rows by clicking the top row and then **SHIFT**+click the bottom row.
2. Right-click the selected rows to display the *Set Import Options* dialog box, select **Do not import**, and then click **Apply**.
3. In the *DNS Zones* table, clear the **Import** checkbox for the default view.
4. Select **corpxyz.com**, **lab.corpxyz.com** and all the corresponding reverse-mapping zones.  
**Tip:** You can use **SHIFT**+click to select multiple contiguous rows and **CTRL**+click to select multiple noncontiguous rows.
5. Right-click the selected rows, and then select **Set Import Options**.
6. In the *Set Import Options* dialog box, enter the following, and then click **Apply**:
  - **Set Zone Type:** No change
  - **Set Import Option:** No change
  - **Set View:** default
  - **Set Member:** HQ-Group master
7. Select **site1.corpxyz.com** and all the reverse-mapping zones with 1 in the second octet in the zone name (1.1.10.in-addr.arpa, 2.1.10.in-addr.arpa, 3.1.10.in-addr.arpa, and so on).
8. Right-click the selected rows, and then select **Set Import Options**.
9. In the *Set Import Options* dialog box, make the same selections as in 6, but choose **Site1-Group** master from the **Set Member** drop-down list.
10. Similarly, select **site2.corpxyz.com** and all the reverse-mapping zones with 2 in the second octet in the zone name.
11. Right-click the selected rows, and then select **Set Import Options**.
12. In the *Set Import Options* dialog box, make the same selections as in 6, but choose **Site2-Group** master from the **Set Member** drop-down list.

## Importing DNS Data

1. Click **Import**.  
The wizard imports the global DNS parameters and zone-specific configuration settings from the named.conf file and performs a zone transfer of the data from the legacy server.
2. Use the *Data Import Wizard Log* to monitor progress and review results afterward.  
The log lists all the zones that the wizard imports and concludes with a total of all the successfully and unsuccessfully imported zones.  
Note if the wizard is unable to import a zone, an error message with an explanation appears in the log.
3. To close the *Data Import Wizard*, click **Exit**. This closes the *Data Import Wizard Log* as well.

## After Using the Wizard

After you import data, you must restart services on the Grid Master and delete the A records for the legacy servers from the corpxyz.com zone. You can also confirm that the imported data is correct and complete by checking the DNS configuration and the forward- and reverse-mapping zones.

1. Log in to the Grid Master (10.0.1.10), select the **Grid** tab, expand the Toolbar, and then click the Restart Services icon.  
Note when importing data through the wizard rather than entering it through the GUI, the Restart Services icon does not change to indicate you must restart service for the appliance to apply the new data. Still, restarting service on the Grid Master is necessary for the imported configuration and data to take effect.

2. To remove A records for the legacy servers, from the **Data Management** tab, select **DNS** tab -> **Zones** tab -> **corpxyz.com**.
3. Expand the **Records** section, select the following A records in the *corpxyz.com* zone, and then click the Delete icon:
  - a. ns1 (for 10.0.1.5)
  - b. ns2 (for 10.0.2.5)
  - c. ns3.site1.corpxyz (for 10.1.1.5)
  - d. ns4.site3.corpxyz (for 10.2.1.5)
4. Remove the respective A records for legacy servers from the *site1.corpxyz* and *site3.corpxyz* subzones.
5. To check the imported DNS configuration file, from the **Data Management** tab, select the **DNS** tab -> **Members** tab -> **10.0.1.10** checkbox. Expand the Toolbar and click **View** -> **View DNS Configuration**. Note that if you do not see the imported DNS configuration file, ensure that you enabled DNS and restarted the services.
6. Scroll through the DNS configuration log to check that each imported zone has an allow-update statement like the following one for the `10.1.10.in-addr.arpa` reverse-mapping zone:

```
zone "10.1.10.in-addr.arpa" in {
    ...
    allow-update { key DHCP_UPDATER; 10.0.2.10; 10.1.1.10; 10.2.1.10; };
    ...
};
```

Enable DHCP and Switch Service to the Grid

Finally, you must enable DHCP service on the three Grid members at 10.0.2.10, 10.1.1.10, and 10.2.1.10, and switch DNS and DHCP service from the legacy DNS and DHCP servers to them.

1. Log in to the Grid Master (10.0.1.10) and from the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **10.0.2.10** checkbox. Expand the Toolbar and click **Start**.
2. Repeat step 1 to enable DHCP on 10.1.1.10 and 10.3.1.10. Note to start the DNS service, as described in [Starting and Stopping the DNS Service](#). The Grid members are ready to serve DHCP and DNS, and send DDNS updates.
3. Take the legacy DHCP and DNS servers offline.

## Managing a Grid

After you configure a Grid Master and add members, you might need to perform the following tasks:

- [Changing Grid Properties](#)
- [Configuring Security Level Banner](#)
- [Configuring Notice and Consent Banner](#)
- [Configuring Informational Level Banner](#)
- [Configuring Recursive Deletions of Networks and Zones](#)
- [Setting the MTU for VPN Tunnels](#)
- [Removing a Grid Member](#)
- [Promoting a Grid Master Candidate](#)
  - [Testing the Connection of the Master Candidate with the Grid Members](#)
  - [Promoting the Master Candidate](#)
  - [Reconnecting Groups After Grid Master Candidate Promotion](#)
- [Enabling Read-only API Access on the Grid Master Candidate](#)

## Changing Grid Properties

You can change a Grid name, its shared secret, and the port number of the VPN tunnels that the Grid uses for communications. Note that changing the VPN port number, time zone, date or time requires a product restart. To modify the properties of a Grid, do the following:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and select **Grid Properties** -> **Edit**.
3. In the *Grid Properties* editor, select the **General** tab -> click the **Basic** tab, and then modify any of the following:
  - **Grid Name:** Type the name of a Grid. The default name is *Infoblox*.
  - **Shared Secret:** Type a shared secret that all Grid members use to authenticate themselves when joining the Grid. The default shared secret is *test*.
  - **Shared Secret Retype:** Type the shared secret again to confirm its accuracy.
  - **Time Zone:** Choose the applicable time zone from the drop-down list.
  - **Date:** Click the calendar icon to select a date or enter the date in YYYY-MM-DD format.
  - **Time:** Click the clock icon to select a time or enter the time in HH:MM:SS format.
  - **VPN Port:** Type the port number that the Grid members use when communicating with the Grid Master through encrypted VPN tunnels. The default port number is 1194. For more information about port numbers for grid communication, see [Creating a Grid Master](#).
  - **Enable Recycle Bin:** Select the checkbox to enable the Recycle Bin. The Recycle Bin stores deleted items when the user deletes Grid, DNS, or DHCP configuration items. Enabling the Recycle Bin allows you to undo deletions and to restore the items on the appliance at a later time. If you do not enable this feature, deleted items from the GUI are permanently removed from the database.
  - **Audit Logging:** Select one of the following:
    - **Detailed:** This is the default type. It is automatically selected. It provides detailed information on all administrative changes such as the date and time stamp of the change, administrator name, changed object name, and the new values of all properties.
    - **Brief:** Provides information on administrative changes such as the date and time stamp of the change, administrator name, and the changed object name. It does not show the new value of the object.
    - **WAPI Detailed:** Select this option to view detailed WAPI (RESTful API) session information logs for successful WAPI calls such as PUT, POST, and DELETE. You can view the URI, InData and response time for each WAPI call. For more information, see [Monitoring Tools](#).
  - In the *Grid Properties* editor, select the **General** tab -> click the **Advanced** tab (or click Toggle Advanced Mode) and modify any of the following:
    - **Enable GUI Redirect from Member:** Select this checkbox to allow the appliance to redirect the Infoblox GUI from a Grid member to the Grid Master.  
Note that if read-only API access is enabled for a Grid Master Candidate, then selecting the **Enable GUI Redirect from Member** checkbox for the Grid Master Candidate does not redirect the Infoblox GUI from the Grid Master Candidate to the Grid Master. For more information about enabling read-only API access on a Grid Master Candidate, see [Enabling Read-only API Access on the Grid Master Candidate](#) below.
    - **Enable GUI/API Access via both MGMT and LAN1/VIP:** Select this checkbox to allow access to the Infoblox GUI and API using both the MGMT and LAN1 ports for standalone appliances and MGMT and VIP ports for an HA pair. This feature is valid only if you have enabled the MGMT port. For information about enabling the MGMT port, see [Appliance Management](#).  
Note that the appliance uses the MGMT port only to redirect the Infoblox GUI from a Grid member to the Grid Master even after you enable the **Enable GUI/API Access via both MGMT and LAN1/VIP** feature.
  - **Show Restart Banner:** Select this checkbox to enable the appliance to display the **Restart Banner** at the top of Grid Manager whenever the appliance notifies you that a service restart is required.
  - **Require Name:** Select this checkbox to prompt the administrator to input the username before performing the service restart. When you select this checkbox, the appliance displays the *Confirm Restart Services* dialog box. Enter the username in the **Name** field and click **Restart Services**. For information about restarting service, see [Restarting Services](#).
4. Save the configuration.

If you changed the VPN port number, time zone, date or time, Grid Manager displays a warning indicating that a product restart is required. Click **Yes** to continue, and then log back in to Grid Manager after the application restarts.

## Configuring Security Level Banner

You can publish a security banner that indicates the security level of the Infoblox Grid. It appears on the header and footer of all pages of Grid Manager. The security level can be Top Secret, Secret, Confidential, Restricted, and Unclassified. Each message type is associated with a predefined security level color. You can modify this color at any point of time. Grid Manager automatically uses an appropriate contrasting text font color that goes with the banner color. Only superusers can configure and enable this feature.

To configure the advanced security level banner for a Grid:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and select **Grid Properties -> Edit**.
3. In the *Grid Properties* editor, select the **Security** tab -> **Advanced** tab.
4. Complete the following:
  - **Enable Security Banner:** Select this to enable the display of the security banner.
  - **Security Level:** From the drop-down list, select the security level for the banner.
  - **Security Level Color:** The default color is displayed in the drop-down. If necessary, using the drop-down list, select the required color for the security level banner. When you change the security level, Grid Manager resets default color for that level.
  - **Classification Message:** Enter the message you want to display in the security banner. You can enter up to 190 characters.
5. Save the configuration.

Security banner that you have configured will appear on the header and footer of the Grid Manager screen including the Login screen.

## Configuring Notice and Consent Banner

You can configure and publish a notice and consent banner as the first login screen that includes specific terms and conditions you want end users to accept before they log in to the Infoblox Grid. When an end user tries to access Grid Manager, this banner is displayed as the first screen. Before accessing the login screen of the Grid Manager, the user must accept the terms and conditions displayed on the consent screen. Only superusers can configure and enable this feature.

To configure the notice and consent banner, do the following:

1. From the **Grid** tab, open the **Grid Manager** tab.
2. Expand the **Toolbar** and select **Grid Properties -> Edit**.
3. In the *Grid Properties* editor, select the **Security** tab -> **Advanced** tab, and then specify the following:
  - **Enable Notice and Consent Banner:** Select this checkbox to enable the display of the notice and consent banner. In the text field, enter the message that you want the banner to show. The message cannot exceed 10,000 characters.
4. Save the configuration.

This banner appears as the first screen when users access Grid Manager. Users must read the terms and conditions and then click **Accept** on the consent screen before they can access the login screen of Grid Manager.

## Configuring Informational Level Banner

You can publish the informational banner for multiple uses, such as to indicate whether the Infoblox Grid is in production or a lab system. The banner can also be used for issuing messages of the day. The informational level banner appears on the header of the Grid Manager screen. You can publish the banner information you want and set the banner color. Grid Manager automatically uses an appropriate contrasting text font color that goes with the banner color. Only superusers can configure and enable this feature.

To configure the advanced informational banner for a Grid, do the following:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and select **Grid Properties -> Edit**.
3. In the *Grid Properties* editor, select the **General** tab -> **Advanced** tab

4. Specify the following:
  - **Enable informational GUI Banner:** Select this checkbox to enable the display of the informational banner message.
  - **Banner Color:** The default color is displayed in the drop-down. If necessary, using the drop-down list, select the required color for the informational level banner.
  - **Message:** Enter the message you want to display in the informational banner. You can enter up to 190 characters.
5. Save the configuration.  
Informational banner appears on the header of the Grid Manager screen.

## Configuring Recursive Deletions of Networks and Zones

Use Grid Manager, to configure a group of users that are allowed to delete or schedule the deletion of a network container, its child objects, a zone and the zone's child objects. For instructions on deleting a network container or a zone, see [Deleting Network Containers](#) and [Removing Zones](#).

When you select **All Users** or **Superusers**, these users can choose to delete a parent object and re-parent its child objects, or they can choose to delete a parent object and all its child objects. These options appear only if a network container or a zone has child objects. For instructions on scheduling recursive deletion of network containers and zones, see [Scheduling Recursive Deletions of Network Containers and Zones](#).

When you select **Nobody**, all the users can delete the parent object only. All the child objects, if any, are re-parented. For more information about scheduling deletions, see as described in [Scheduling Deletions](#). Note that you can restrict specific users to perform recursive deletions of network containers and zones only through Grid Manager. These settings do not prevent other users from performing recursive deletions through the API.



### Note

You must have Read/Write permission to all the child objects in order to delete a parent object. Recursive deletion is applicable to all zone types except stub and forward-mapping zones.

The appliance puts all deleted objects in the Recycle Bin, if enabled. You can restore the objects if necessary. When you restore a parent object from the Recycle Bin, all its contents, if any, are re-parented to the restored parent object. For information about Recycle Bin, see [Finding and Restoring Data](#).

To configure the group of users to perform recursive deletions:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and select **Grid Properties -> Edit**.
3. In the *Grid Properties* editor, select the **General** tab -> **Advanced** tab.
4. Under **Present the option of recursive deletion of networks or zones to**, select one of the following:
  - **All Users:** Select this to allow all users, including superusers and limited-access users, to choose whether they want to delete the parent object and its contents or the parent object only when they delete a network container/network or a zone. This option is selected by default.
  - **Superuser:** Select this to allow only superusers to choose whether they want to delete the parent object and its contents or the parent object only when they delete a network container/network or a zone.
  - **Nobody:** When you select this, users can only delete the parent object (network container or zone). All child objects, if any, are re-parented.
5. Save the configuration.

## Setting the MTU for VPN Tunnels

You can configure the VPN MTU (maximum transmission unit) for any appliance with a network link that does not support the default MTU size (1500 bytes) and that cannot join a Grid because of this limitation. If an appliance on such a link attempts to establish a VPN tunnel with a Grid Master to join a Grid, the appliance will receive a PATH-MTU error, that indicates that the path MTU discovery process has failed. For information about the MTU discovery process, see *RFC 1191, Path MTU Discovery*.

To avoid this problem, set a VPN MTU value on the Grid Master for any appliance that cannot link to it using a 1500-byte MTU. When the appliance contacts the master during the key exchange handshake that occurs during the Grid-joining operation, the master sends the appliance the MTU setting to use.

To set the VPN MTU for a Grid member, do the following:

1. In the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox -> Edit icon.
2. Select the **Network** -> **Advanced** tab of the *Grid Member Properties* editor.
3. In the **VPN MTU** field, enter a value between 600 and 1500.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Removing a Grid Member

You might want or need to remove a member from a Grid, perhaps to disable it or to make it an independent appliance or an independent HA pair. Before you remove a member, make sure that it is not assigned to serve any zones or networks. To remove a Grid member do the following:

1. In the **Grid** tab, open the **Grid Manager** tab -> **Members** tab
2. Select the *Grid\_member* checkbox
3. Click the Delete icon.

## Promoting a Grid Master Candidate

To promote a Grid Master candidate to a Grid Master, you must have already designated a member as a Grid Master Candidate, by selecting the **Master Candidate** option in the **General** tab of the *Grid Member Properties* editor. You can designate any member as a Grid Master Candidate. The Grid Master Candidate gets a complete copy of the Grid database. Therefore, Infoblox recommends that you configure the same appliance models for the Grid Master and Grid Master Candidates. By default, the Grid Master promotion uses UDP port 1194. Make sure that the UDP 2114 and UDP 1194 ports are open between the Grid members and a newly designated Grid Master. During a Grid Master promotion, the newly promoted Grid Master continuously contacts all Grid members, including the original Grid Master on the UDP port 2114, until it reaches them. Upon reaching them, the newly promoted Grid Master notifies all Grid members that it is the new Grid Master. Next, the Grid Members restart and attempt to establish normal Grid communications (via BloxSync) with the newly promoted Grid Master. Before promoting a Grid Master Candidate, check your firewall rules to ensure that the Master Candidate can communicate with all the Grid members. For information about grid communications, see [About Grids](#).



### Note

Before promoting a Grid Master Candidate, ensure that valid client SSL certificates are installed. For more information about installing certificates, see [Managing Certificates](#).

## Testing the Connection of the Master Candidate with the Grid Members

Before promoting a Grid Master Candidate, check whether the Grid Master Candidate is connected to the rest of the Grid members, by scheduling a test promotion. You can do this either by using Grid Manager or by using the NIOS CLI. For information about scheduling a test promotion by using the NIOS CLI, see [show test\\_promote\\_master](#) and [set test\\_promote\\_master](#).

The connection of the Grid Master Candidate to the rest of the Grid members is checked by sending specifically crafted test packets from the Grid Master Candidate and checking whether the Grid members are able to receive these packets.

To test the connection of the Grid Master Candidate with the Grid members, complete the following:

1. In the **Grid** tab -> **Grid Manager** tab, expand the Toolbar, and then click **GMC Promote Test**.
2. In the *GMC Promote Test* editor, do the following:
  - a. Click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, do one of the following:
    - To run a test promotion immediately, select **Now**.
    - To schedule a test promotion to run later, select **Later**, and enter a date, time, and time zone.
  - b. From the **Select GMC** drop-down list, select the Grid Master Candidate that you want to promote to Grid Master.
  - c. In the **Timeout (secs)** field, set the timeout for the packet to be received in seconds. That is, if the packet is not received by the Grid members within this timeout, the connection is deemed to have failed.
  - d. Select the **Continuous Testing** checkbox if you want the Grid Master Candidate to send packets to the selected Grid members on a continual basis. The maximum period of time for which packets can be sent is 120 seconds.



- e. In the **Members** table, select the Grid members to which the Grid Master Candidate must establish a connection.
3. Click **Start** to start the test promotion. You can click **Stop** at any time to stop the test promotion.
4. Click **GMC Promotion Test Results** to view the status of the test promotion.



#### Notes

- You cannot upgrade the Grid during a test promotion.
- You can test promote of only one Grid Master Candidate at a time.
- If new members are added when a test promotion is in progress, connection between the new members and the Grid Master Candidate will not be tested.
- If Threat Protection is enabled in the Grid and the member running the Threat Protection service is in the list of tested members, you must set the value in the **Timeout** field to at least 30 seconds. This is because Threat Protection needs to publish a new rule that allows traffic to pass from tested members. If you set a lower timeout value, the packets may be dropped, and the test will report that the member cannot connect to the tested Grid Master Candidate.
- Communication between DUT and Grid Master is not tested because of firewall complications and running the OpenVPN connection. Communication is supposed to be already checked and DUT is already connected to Grid Master.
- You cannot run continuous testing when a regular test is in progress and you cannot run a regular test when continuous testing is in progress.
- If multiple public cloud instances such as AWS, Azure, GCP and so on are configured as the Grid Master Candidate, ensure that these instances are able to communicate with other public cloud instances. Otherwise, the Grid Master Candidate promote test does not work.
- When you configure a Grid Master Candidate which includes an External NTP server and when you promote a Grid Master Candidate to Grid Master, the External NTP is enabled in the Grid Master Candidate. In case you try to edit the member properties an error message is displayed. Therefore, Infoblox recommends that you remove the External NTP configuration before you promote the Grid Master Candidate.

## Promoting the Master Candidate

To promote a Master Candidate, you can make a direct serial connection to the console port on the active node of an HA Candidate or to the console port on a single Candidate. You can also make a remote serial connection (using SSH v2) to the candidate. Enter the following Infoblox CLI command to promote a Master Candidate:

```
set promote_master.
```

You can do one of the following to promote a Master Candidate:

- Immediately notify all Grid members about the promotion.
- Set a sequential notification to provide wait time for Grid members to join the new Grid Master. Staggering the restarts of Grid members can minimize DNS outages. The sequential order for Grid members to join the new Grid Master begins with the old Grid Master and then the Grid members in FQDN order. The default delay time is 120 seconds. You can configure the delay time from a minimum of 30 seconds up to 600 seconds.



#### Notes

- During a Grid Master promotion, ensure that you do not designate a Grid member as a Grid Master Candidate or promote a Master Candidate. In addition, wait up to two hours since the last promotion to perform another Grid Master promotion. Otherwise, you might experience unnecessary member reboots. Whenever possible, separate any operations that require product restarts by at least an hour.
- When a Grid Master Candidate is selected as a subscribing member, then after Grid Master Candidate promotion, the subscription still takes place through the previous Grid Master Candidate member which is now a Grid member.



To promote a Grid Master Candidate, do the following:

1. Establish a serial connection (through a serial console or remote access using SSH) to the Master Candidate. For information about making a serial connection, as described in Method 2-Using the CLI, see [Deploying a Single Independent Appliance](#).
2. At the CLI prompt, use the command `set promote_master` to promote the Master Candidate and send notifications to all Grid members immediately, or promote the Master Candidate to the Grid Master immediately and specify the delay time for the Grid members to join the new Grid Master. For more information about the command, refer to the *Infoblox CLI Guide*.
3. To verify the new master is operating properly, log in to the Infoblox Grid Manager on the new master using the VIP address for an HA master or the IP address of the LAN1 port for a single master.
4. Check the icons in the **Status** column. Also, select the master, and then click the Detailed Status icon in the table toolbar. You can also check the status icons of the Grid members to verify that all Grid members have connected to the new master. If you have configured delay time for Grid member notification, it will take some time for some members to connect to the new master. You can also check your firewall rules and log in to the CLI to investigate those members.



Note that when you promote the Master Candidate to a Grid Master, the IP address will change accordingly. If you have configured a FireEye appliance, then any changes in the Grid Master IP address, FireEye zone name, associated network view or the DNS view will affect the **Server URL** that is generated for a FireEye appliance. The FireEye appliance will not be able to send alerts to the updated URL when there is a change in the IP address. You must update the URL in the FireEye appliance to send alerts to the NIOS appliance. For more information, see [Configuring FireEye RPZs](#).

## Reconnecting Groups After Grid Master Candidate Promotion

This feature gives you more control over the Grid Master Candidate promotion, minimizes service outages by allowing you to group the members and schedule a time for the groups to reconnect to the newly promoted Grid Master. As soon as the scheduled time arrives, members of Grid Master Candidate groups will re-connect to the newly promoted master.

To schedule a group reconnection to the newly promoted Grid Master Candidate, do the following:

1. From the Grid tab -> **Grid Manager** tab, expand the Toolbar, and then click **GMC Group Promotion**.
2. In the **GMC Group Promotion Schedule** editor, specify the following:
  - **Activate GMC Group Promotion Schedule:** Select this option to enable the scheduled reconnection of the group after the Grid Master Candidate is promoted.
  - Click the + icon and specify the following in **Add GMC Group Wizard**:
    - **Name:** Provide the group's name.
    - **Promotion Policy:** Select either **Simultaneously** or **Sequentially**, as required.
      - Simultaneously:** Select this option to simultaneously reconnect the group members after the Grid Master Candidate promotion at the same time.
      - Sequentially:** Select this option to sequentially reconnect the group members after Grid Master Candidate promotion in a sequence. Note that when you select sequentially, each group member joins the Grid master in a sequence with an interval of 30 seconds.
    - **Time Zone:** Select a time zone that applies to the start time you enter. If this time zone is different from the Grid time zone, the appliance converts the time you enter here based on the Grid time zone, after you save this schedule. When you display this schedule again, it displays the converted time. Selecting the time zone here does not affect any time zone settings in the Grid. (For information about setting the Grid and member time zones, see [Managing Time Settings](#)). After the Grid Master Candidate promotion, members will reconnect based on the selected time zone.
    - **Date:** Enter a start date of the group members reconnecting after Grid Master Candidate promotion in YYYY-MM-DD (year-month-day) format. You can click the calendar icon to select a date from the calendar widget.
    - **Time:** Enter a start time of the group members reconnecting after Grid Master Candidate promotion in hh:mm:ss AM/PM (hour:minute:second in AM or PM) format. You can select a time from the drop-down list.

- **Comment:** Enter your comments.
  - Click **Next**.
  - In the **Members Assignment** wizard, select the Grid member(s) to add to the newly created group.
3. **Save** and **close** the wizard.

To modify an existing group, on the **GMC Group Promotion Schedule** editor:

1. Click Edit icon, and modify the changes in **Add GMC Group Wizard**.
2. **Save** and **close** the wizard.

To delete an existing Grid Master Candidate group, do the following in the **GMC Group Promotion Schedule** editor:

1. Click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.

After enabling the Grid Master Candidate group promotion, use the `set promote_master` CLI command to start the Grid Master Candidate promotion.

Use the `set gmc_promotion disable` CLI command to disable the **Activate GMC Group Promotion Schedule** option. Note that, this command can be executed on Grid Master and Grid Master Candidate. For more information see, [set gmc\\_promotion](#).



#### Notes

- If the **Activate GMC Group Promotion Schedule** option is not enabled on Grid Master, and if you choose to continue with Grid Master Candidate promotion using the command, `set promote_master`, then the Grid Master Candidate promotion works as described in [Promoting the Master Candidate](#).
- By default, all the members are part of the **Default** group. The Grid Master Candidate group members can be customized according to your requirement; however Grid Master Candidate cannot be customized as Grid Master Candidate is part of the **Default** group only.
- If you want to reconnect members of any group to the newly promoted Grid Master, irrespective of the scheduled time, you can click **Join Group Now** option, by selecting the following **Join Group Now** icon,



in **GMC Group Promotion Schedule** editor. This works only during the promotion of a Grid Master Candidate. That is, **Join Group Now** is activated (enabled) only during the promotion of Grid Master Candidate group; it is disabled after the scheduled time of all the groups expires after the Grid Master Candidate is promoted. For any offline member, the **Join Group Now** will be disabled 8 hours after the Grid Master Candidate promotion.

- The **Add GMC Group Wizard** in the **GMC Group Promotion Schedule** editor is available only for future schedules. The maximum scheduled time for the promotion of any Grid Master Candidate group is 8 hours.
- We do not recommend enabling a schedule Grid upgrade and **GMC Group Promotion Schedule** at the same time.
- The **Time Zone** for any group, displays the Grid Manager's time zone, and if there are any member(s) in the group, the **Time Zone** automatically reflects the first group member(s) time zone.
- The scheduled **Time** displays the new time zone, if the **Time Zone** is modified or if a member is moved across different groups.
- During the Grid Master Candidate promotion, if a Grid member is offline, Grid Manager continuously attempts to connect to the offline Grid member for every 60 seconds.
- If the **GMC Group Promotion Schedule** editor is disabled after Grid upgrade, then you can unset the previously triggered Grid Master Candidate promotion, by using the CLI command `set gmc_promotion forced_end`. It is recommended to run this command when the Grid is completely upgraded. For more information see, [set gmc\\_promotion](#).

## Enabling Read-only API Access on the Grid Master Candidate

You can enable read-only API access on the Grid Master Candidate to provide additional scalability of read/write API requests on the Grid Master, which in turn improves the performance of the Grid Master. The read-only API access is disabled by default for new installations.

When you enable read-only API access on an HA Grid Master Candidate, you can access the API service only on an active node. If the API service is disabled for an admin group, the users in the admin group cannot access read-only API service on the Grid Master Candidate, even though read-only API access is enabled for the Grid Master Candidate. Also, the users in the admin group should have at least read-only permission to access the API service.



### Note

When you upgrade the Grid Master Candidate to NIOS 7.1 and later, read-only API access is disabled. But when you upgrade the Grid Master Candidate from NIOS 7.1 to a later release with read-only API access enabled, then this setting is retained after the upgrade is completed.

The appliance logs all API logins in the audit log and syslog. You can view the audit log and syslog of the Grid Master Candidate under the **Administration** -> **Logs** tab.

To enable read-only API access on the Grid Master Candidate:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_Master\_Candidate* checkbox, and then click the Edit icon.
  - In the *Grid Member Properties* editor, select the **General** tab -> **Basic** tab, and then do the following:
    - Read Only API access:** This field is displayed only when the Grid member is designated as a Master Candidate. Select this checkbox to enable read-only API access on the Grid Master Candidate. Enabling this checkbox will only allow read-only API access and not write API access. Note that if you enable this checkbox, you cannot access the GUI using the IP address of the Grid Master Candidate.
2. Save the configuration.

## Grid Bandwidth Considerations

Infoblox Grid technology relies upon database replication for its core functionality. When designing a Grid, it is important to consider the amount of traffic generated by this replication and the overall number of Grid members. Other communication between Grid members (such as log retrieval and monitoring functions) occurs as well. All of this traffic is securely communicated between the Grid Master and Grid members through encrypted VPN tunnels.

One component of the traffic through the tunnels is database replication traffic. There are three types to consider:

- **Complete database replication to a Master Candidate** — Occurs when a Master Candidate joins or rejoins a Grid. The Grid Master sends the complete database to a Master Candidate so that it has all the data it needs if it ever becomes promoted from member to master.
- **Partial database replication** — Occurs when an appliance or HA pair joins or rejoins the Grid as a regular member (which is not configured as a Master Candidate). The Grid Master sends it the section of the database that mainly applies just to the member.
- **Ongoing database updates** — Occurs as changes are made to the Grid configuration and data. The Grid Master sends all ongoing database updates to Master Candidates and individual member-specific updates to regular members.

If there are no or very few DNS dynamic updates, and no or very few DHCP lease offers and renewals issued, then this type of replication traffic is minimal.

If there are many DDNS (dynamic DNS) updates (many per second) and/or many DHCP lease offers and renewals (many per second), then the replication traffic is the largest component of the VPN traffic among Grid members.

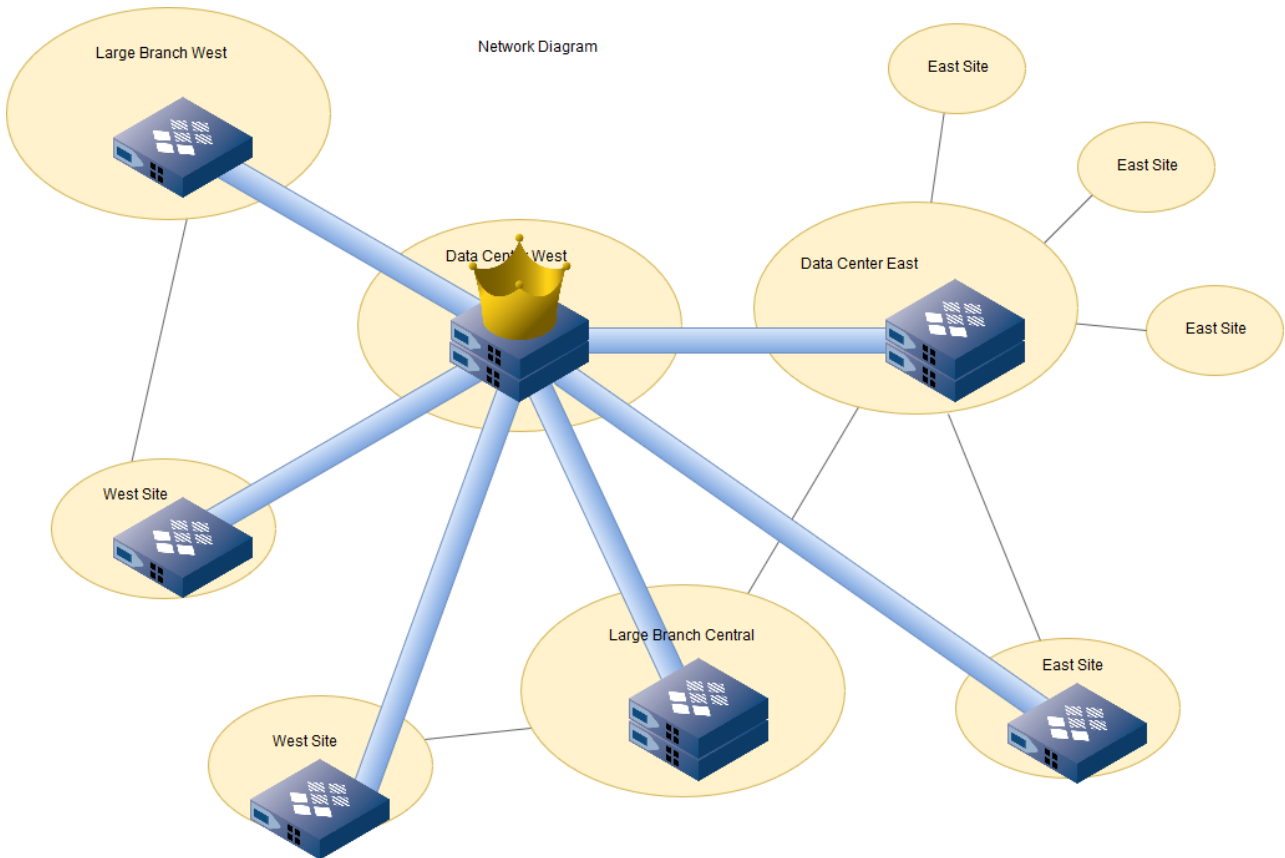


### Note

A Grid Master replicates data to single members and to the active node of HA members. The active node then replicates the data to the passive node in the HA pair.

At a minimum, there must be 256 Kbps (kilobits per second) bandwidth between the Grid Master and each member, with a maximum round-trip delay of 500 milliseconds. For ongoing database updates, the amount of data sent or received is 15 Kb for every DDNS update, and 10 Kb for every DHCP lease -offer/renew. The baseline amount for heartbeat and other maintenance traffic for each member is 2 Kbps. Measure the peak DNS and DHCP traffic you see in your network to determine the bandwidth needed between the Grid Master and its members for this activity. For example, you might decide to place your Grid members in the locations shown in the below figure.

### Grid Deployment



In this example, the Grid Master is optimally placed in the Data Center West. There are a total of seven members: the HA Grid Master, three HA members, and three single members. If all the members are Master Candidates, the Grid Master replicates all changes to the other six members. Assuming that the master receives 20 dynamic updates per minute and 40 DHCP lease renews per minute, the calculation for Grid bandwidth is:

$20 \text{ DDNS updates/minute} / 60 \text{ secs} = 0.333 \text{ DDNS updates/sec} * 15 \text{ Kb} = 5 \text{ Kbps} * 6 \text{ members} = 30 \text{ Kbps}$

$40 \text{ DHCP leases/minute} / 60 \text{ secs} = 0.666 \text{ DHCP leases/sec} * 10 \text{ Kb} = 6.7 \text{ Kbps} * 6 \text{ members} = 40.2 \text{ Kbps}$

$2 \text{ Kbps of Grid maintenance traffic} * 6 \text{ members} = 12 \text{ Kbps}$

Total 82.2 Kbps

Another component is the upgrade process. See [Upgrading NIOS Software](#) for more information. Bandwidth requirements, database size, and update rate determine the maximum size of the Grid you can deploy. Based on the various factors discussed above, you can determine the amount of bandwidth your Grid needs. If your calculations exceed the available bandwidth, then you might need to modify your deployment strategy, perhaps by splitting one large Grid into two or more smaller ones.



#### Note

This calculation does not take into account existing traffic other than DNS and DHCP services, so factor and adjust accordingly.

For international networks, because of bandwidth and delay requirements, a geographical grouping of Grid members might be the best approach. For example, if you have a global presence, it may make the most sense to have a North American Grid, a South American Grid, a European Grid, and an Asia/Pacific Grid.

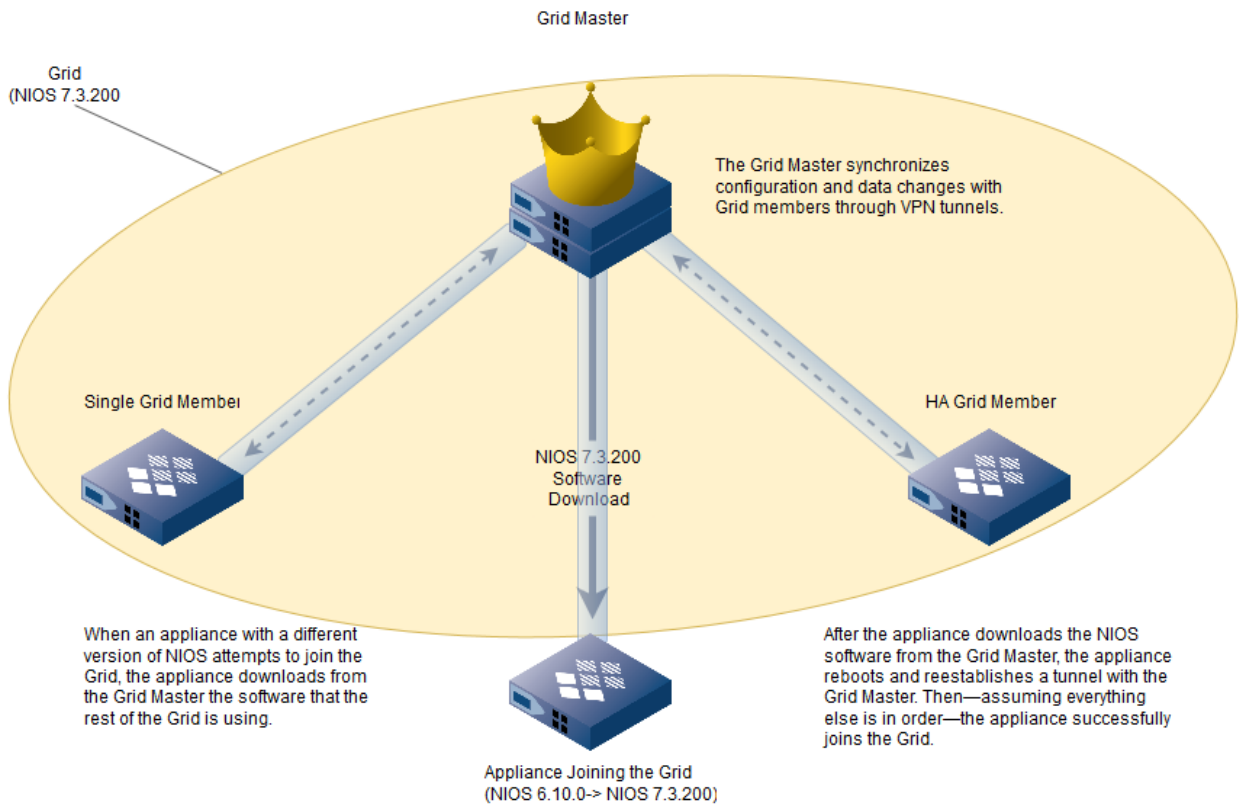
## Automatic Software Version Coordination

When you add an appliance or HA pair to a Grid as a new member, it is important that it is running the same version of software as the other members in the Grid. Infoblox provides two methods for coordinating the software version:

- **Manual Upgrade and Downgrade:** Before adding an appliance or HA pair to a Grid, you can manually upgrade or downgrade the software on the appliance or HA pair to the version used by the rest of the Grid.
- **Automatic Upgrade and Downgrade:** The Grid Master automatically compares the software version of each appliance attempting to enter a Grid with that in use by the rest of Grid. If the versions do not match, the Grid Master downloads the correct version to the new appliance or HA pair.

Note that the Grid Master checks the software version every time an appliance or HA pair joins the Grid. The software version check occurs during the initial join operation and when a member goes offline and then rejoins the Grid.

### *Automatic Upgrade of An Appliance Joining a Grid*



When a single appliance attempts to join the Grid for the first time, the following series of events takes place:

1. The appliance establishes an encrypted VPN tunnel with the Grid Master.
2. The master detects that the software version on the appliance is different from that in the rest of the Grid. For example, the appliance is running NIOS 6.10.0 software but the rest of the Grid is running NIOS 7.3.200 software.
3. The appliance downloads the NIOS 7.3.200 software from the Grid Master.
4. After the upgrade is complete, the NIOS application automatically restarts.
5. After the appliance reboots, it again contacts the Grid Master and step 1 is repeated. Because the software versions now match, the appliance can complete its attempt to join the Grid.

When an HA pair attempts to join the Grid for the first time, the following series of events takes place:

1. The active node of the HA pair establishes an encrypted VPN tunnel with the Grid Master.
2. The master detects that the software version on the node is different from that in the rest of the Grid. For example, the active node is running NIOS 6.10.0 software but the rest of the Grid is running NIOS 7.3.200 software.
3. The appliance downloads the NIOS 7.3.200 software from the Grid Master.
4. After the upgrade is complete, the NIOS application on the active node automatically restarts. This causes an HA failover.
5. The new active node (which was previously the passive node) attempts to join the Grid, repeating steps 1 – 4.
6. When the NIOS application on the currently active node restarts, there is another failover, and the currently passive node becomes active again.
7. The active node again contacts the Grid Master and step 1 is repeated. Because the software versions now match, it can complete its attempt to join the Grid.

## NAT Groups

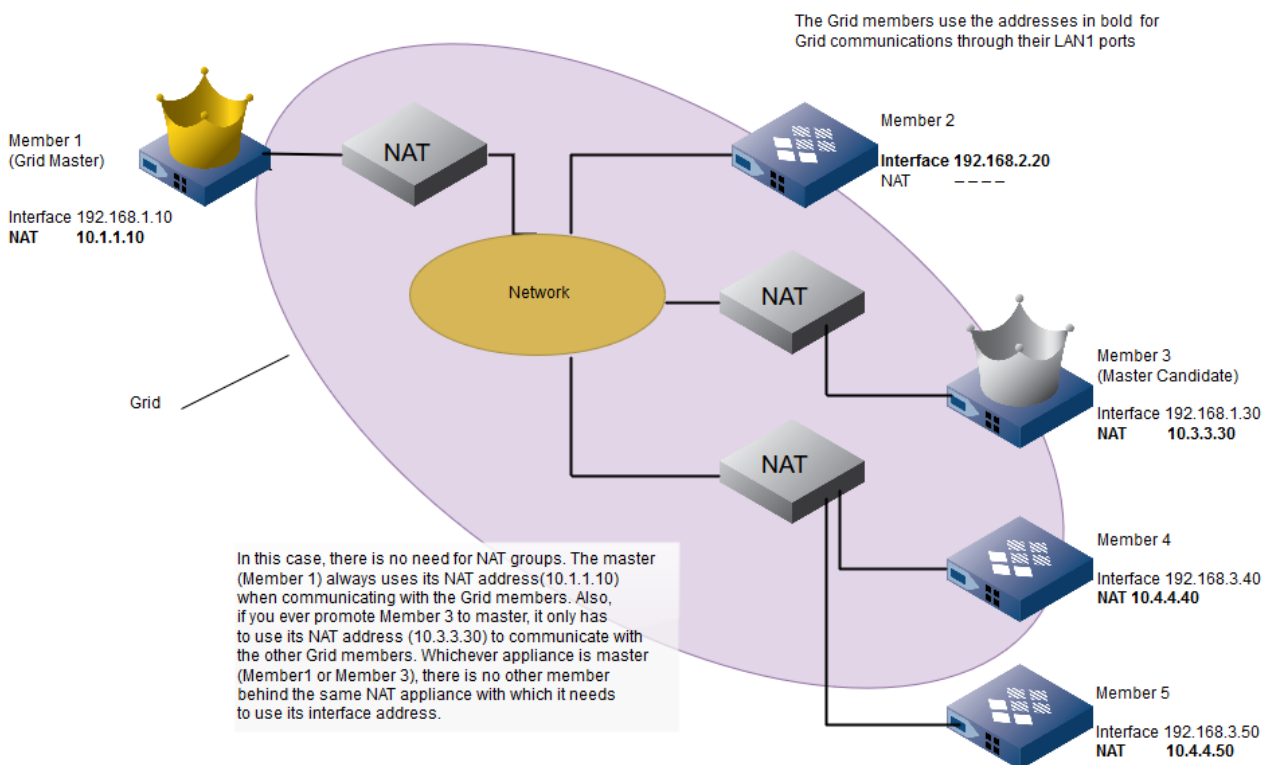


### Note

Infoblox NAT and NAT groups do not support NAT IPv6 operation.

NAT groups are necessary if the Grid Master is behind a NAT appliance and there are members on both sides of that NAT appliance. Any members on the same side as the master go into the same NAT group as the master and use their interface addresses for Grid communications with each other. Grid members on the other side of that NAT appliance do not go in the same NAT group as the master and use the master's NAT address for Grid communications. These other members outside the NAT appliance can—but do not always need to be—in a different NAT group. To see when NAT groups become necessary for Grid communications, compare the figure NAT without NAT Groups below with the figures Grid Master in NAT Group and Grid Master and Master Candidate in NAT Groups.

### *NAT without NAT Groups*

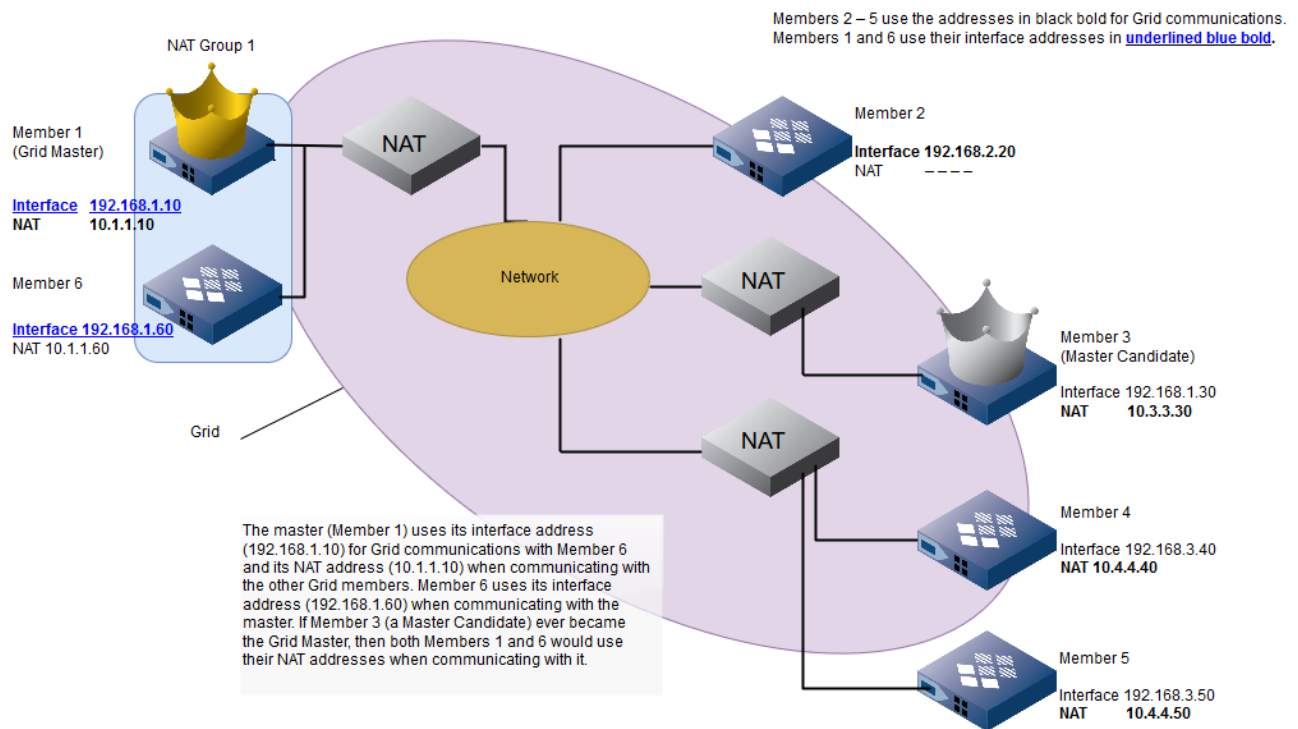


### Note

A single or HA member using its MGMT port for Grid communications cannot be separated from the Grid Master behind a NAT appliance. For more information, see [Using the MGMT Port](#).



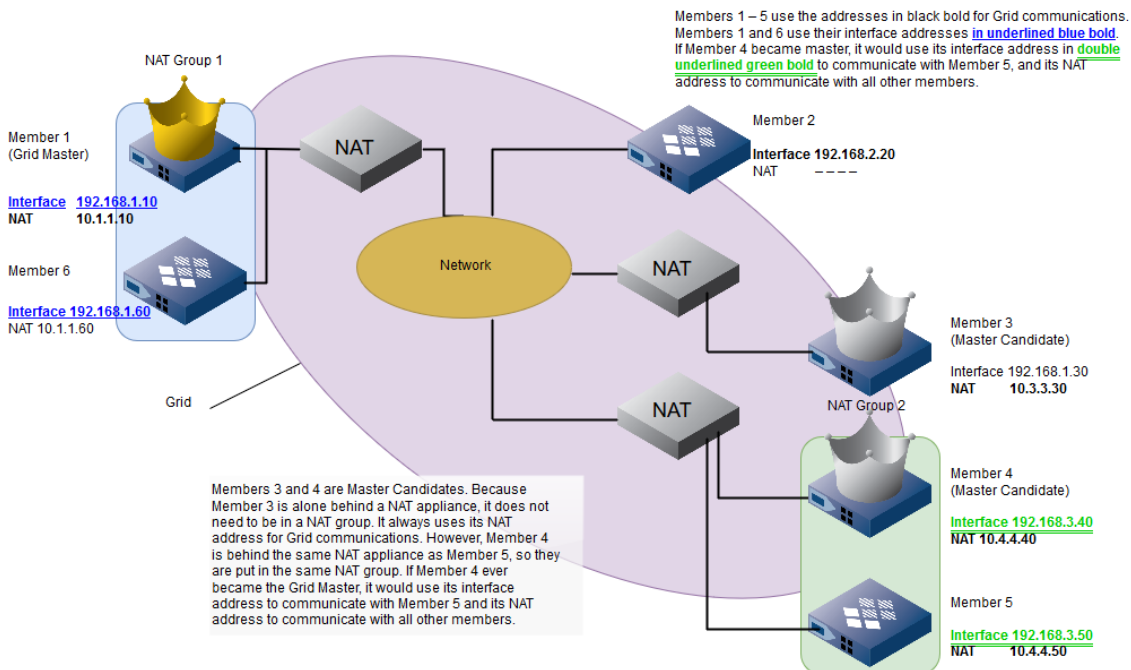
## Grid Master in NAT Group



The same use of NAT groups that applies to a Grid Master also applies to Master Candidates. If there are no other members behind the same NAT appliance as a Master Candidate, then the Master Candidate does not need to be in a NAT group. It always uses its NAT address for Grid communications. If another member is behind the same NAT appliance as the Master Candidate, then both the candidate and that member need to be in the same NAT group so that—if the candidate becomes master—they can use their interface addresses to communicate with each other (see the figure below).

## Grid Master and Master Candidate in NAT Groups





Although some members might not need to be in a NAT group, it is good practice to put all members in NAT groups in anticipation of adding or rearranging Grid members within the network. For example, in NAT without NAT Groups – Grid Master and Master Candidate in NAT Groups, Member 4 did not need to be in a NAT group until it became configured as a Master Candidate. At that point, because Member 5 is also behind the same NAT appliance, it became necessary to create NAT Group 2 and add Members 4 and 5 to it. Similarly, if you add another member behind the NAT appliance in front of Member 3, then you must create a new NAT group and add Member 3 and the new member to it. Always using NAT groups can simplify such changes to the Grid and ensure that NAT appliances never interrupt Grid communications. To create a NAT group:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and select **Grid Properties -> Edit**.
3. In the *Grid Properties* editor, select the **NAT Groups** tab.
4. Click the Add icon, and enter a name in the **Name** field and optionally, a comment in the Comment field.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

To add members to the NAT group:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Select a Grid member and click the Edit icon.
3. In the *Grid Member Properties* editor, select the **Network -> Advanced** tab and complete the following:
  - **Enable NAT Compatibility (IPv4 only):** Select this checkbox. NAT group is not supported by IPv6 appliances.
  - **NAT Group:** From the drop-down list, select the NAT group you previously created.
  - **NAT Addresses:** For a single Grid Master or member, enter the address configured on the NAT appliance that maps to the interface address of the LAN1 port. A single master or member that serves DNS uses this NAT address for Grid communications and—if it serves DNS—DNS messages.

For an HA Grid Master or member, enter the address configured on the NAT appliance that maps to its VIP address. An HA master uses its VIP NAT address when communicating with Grid members. An HA member that serves DNS uses its VIP NAT address for its DNS messages. It uses its LAN1 port NAT address for Grid communications.

- **Node 1 (if HA)**

- **NAT IP Address:** Enter the address configured on the NAT appliance that maps to the interface address of the LAN1 port on Node 1. When Node 1 of an HA member is active, it uses its NAT address for Grid communications.
- **Node 2 (if HA)**
  - **NAT IP Address:** Enter the address configured on the NAT appliance that maps to the interface address of the LAN1 port on Node 2. When Node 2 of an HA member is active, it uses its NAT address for Grid communications.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Deploying Independent Appliances

This section explains how to deploy single independent appliances and independent HA pairs. Independent appliances run NIOS without the Grid upgrade and are deployed independently from a Grid. This chapter includes the following topics:

- [Independent Deployment Overview](#)
- [Deploying a Single Independent Appliance](#)
- [Configuration Example: Deploying a NIOS Appliance as a Primary DNS Server](#)
- [Deploying an Independent HA Pair](#)
- [Configuration Example: Configuring an HA Pair for Internal DNS and DHCP Services](#)
- [Verifying the Deployment](#)
- [Infoblox Tools for Migrating Bulk Data](#)

## Independent Deployment Overview

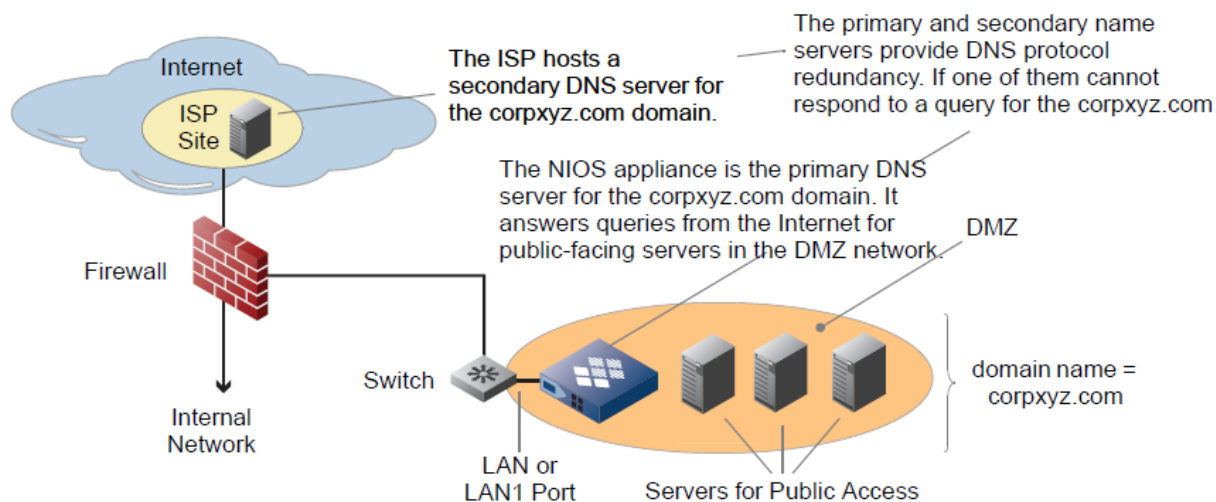


### Note

Infoblox appliances support IPv4 and IPv6 networking configurations in most deployments cited in this chapter. You can set the LAN1 port to an IPv6 address and use that address to access the NIOS UI and the NIOS Setup Wizard. All HA operations can be applied across IPv6. You can also set a dual mode appliance by configuring both IPv4 and IPv6 address for the LAN1 port. Topics in this and following chapters generally use IPv4 examples. Also note that LAN2 and the MGMT port also support IPv6. DNS services are fully supported in IPv6 for the LAN1, LAN2, MGMT and VLAN ports. DHCP services are fully supported in IPv6 for the LAN1 and LAN2 ports. Example networks throughout this chapter use IPv4 addressing.

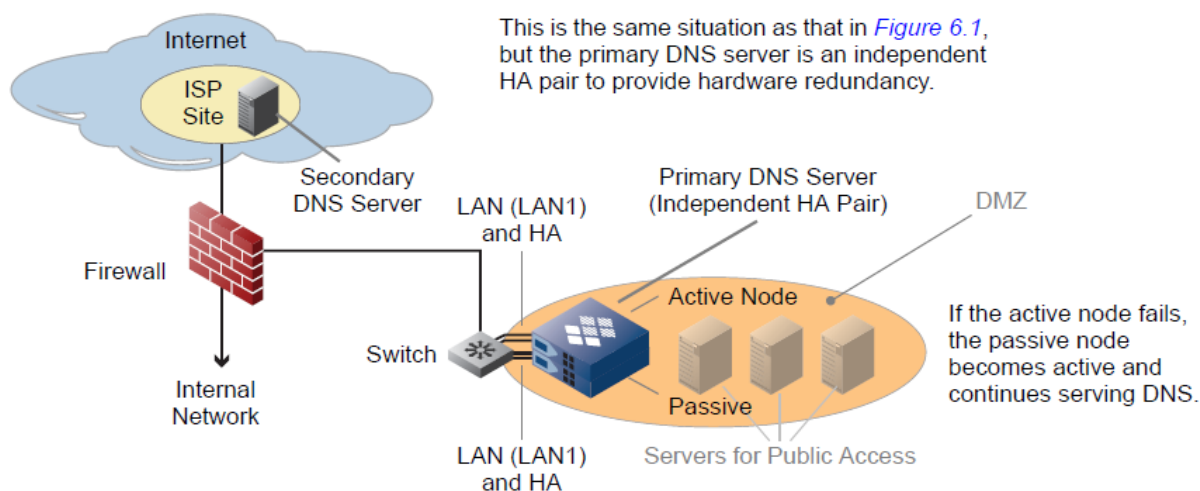
You can deploy the NIOS appliance as a Grid member in an Infoblox Grid or independently as a standalone deployment. NIOS appliances support both IPv4 and IPv6 networks and you can deploy them in either IPv4, IPv6, or dual mode (IPv4 and IPv6). Grids offer many advantages for large organizations while independent deployments can be sufficient for smaller sites. For example, if your ISP hosts one name server to respond to external DNS queries, you can deploy a single independent NIOS appliance as the other name server, as shown in the below figure.

*Single Independent Appliance as a DNS Server*



Using primary and secondary name servers provides DNS protocol redundancy, and configuring two DHCP servers as DHCP failover peers provides DHCP protocol redundancy. However, you can only have hardware redundancy if you deploy appliances in an HA (high availability) pair. Should the active node in an HA pair fail, the passive node becomes active and begins serving data, as shown in the below figure. For more information about HA pairs, see [About HA Pairs](#).

#### Independent HA Pair



#### System Manager GUI

When you deploy an independent appliance, you use System Manager to manage the appliance. Though other chapters in this guide contain information that assumes a Grid deployment and describes the Grid Manager GUI, most of the configuration procedures are applicable to an independent appliance, with the following differences:

- In the Dashboard, there is no *Grid Status* widget, and the *Members Status* widget in Grid Manager is the *System Status* widget in System Manager.
- Functions related to a Grid, such as joining a Grid and managing Grid licenses, do not exist in System Manager.
- The Grid related tabs and functions in Grid Manager are the system related tabs and functions in System Manager.

- Functions related to the **Members** tab in Grid Manager appear in the **Nodes** tab or the Toolbar of another subtab in System Manager.

For example, the following navigation path for a Grid:

- From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox, and then click **HTTPS Cert** -> **Download Certificate** from the Toolbar.

is the following for an independent appliance:

- From the **System** tab, select the **System Manager** tab -> **Nodes** tab, and then click **HTTPS Cert** -> **Download Certificate** from the Toolbar.

## Deploying a Single Independent Appliance

To deploy a single independent NIOS appliance, connect its LAN1 port to the network and change its default IP settings so that it can connect to its surrounding IP address space. The default LAN settings are as follows:

- IP address: 192.168.1.2
- Netmask: 255.255.255.0
- Gateway: 192.168.1.1

When deploying a single independent appliance, you can use one of the following methods to set up the initial configuration:

- *Method 1 - Using the LCD*
  - Requirements: Physical access to a powered up NIOS appliance.
  - Advantage: You do not need any other equipment.
- *Method 2 - Using the CLI*
  - Requirements: A serial connection from your management system to the console port on the NIOS appliance. You can also enable remote console access so that you can use the CLI over a network connection. For more information, see [Restricting GUI/API Access](#).
  - Advantage: You do not need to change the IP address of the management system to connect to the NIOS appliance.
- *Method 3 - Using the Infoblox NIOS Startup Wizard*
  - Requirements: An HTTPS connection from your management system to the LAN1 port on the NIOS appliance.
  - Advantage: The wizard provides step-by-step guidance for changing not only the IP settings for the LAN1 port, but also changing the appliance host name and admin password, setting the system clock, and—if using NTP (Network Time Protocol)—enabling the NIOS appliance to be an NTP client.


Note that you can configure network settings using the Startup wizard any time after you have configured the appliance. To start the wizard, from System Manager, select the **System** tab, and then click **System Properties** -> **Startup Wizard** from the Toolbar.

After you configure the network settings on a single independent appliance, you can migrate data from legacy DNS and DHCP servers to the NIOS appliance. Several tools and methods are available for migrating data and configuration settings. For a list of the available options, see [Infoblox Tools for Migrating Bulk Data](#).

### Method 1 – Using the LCD

Some of the NIOS appliances have an LCD and navigation buttons on the front panel that allow you to view system status and license information, as well as configure network settings for the LAN1 port.

#### *Infoblox LCD and Navigation Buttons*



The LCD panel is on the front of some of the NIOS appliances.

You can deploy a single independent NIOS appliance by setting its LAN1 port IP address, netmask, and gateway through the LCD. This is the simplest method because you do not need anything other than a physical access to the appliance to complete the initial configuration.



**Note**

LCD does not support IPv6 addressing.

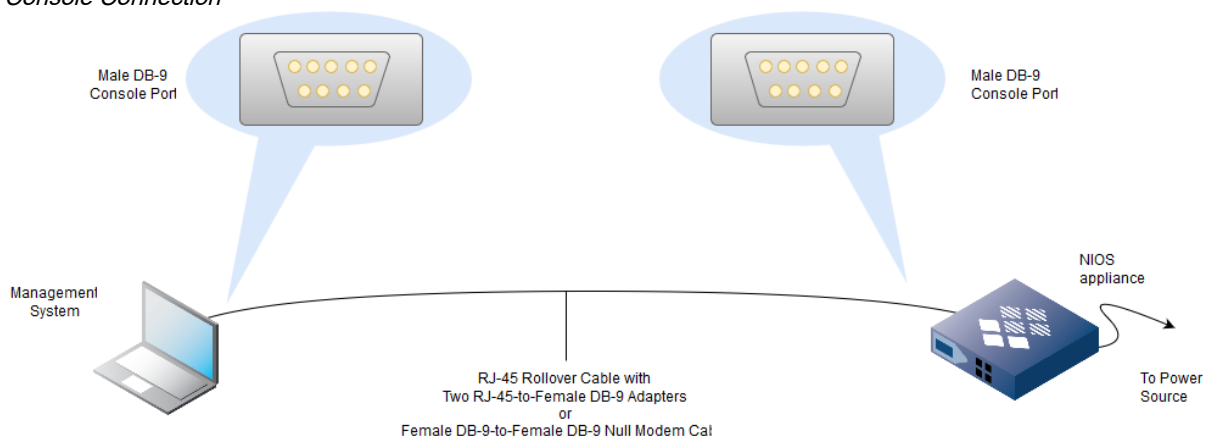
1. Connect the power cable from the NIOS appliance to a power source and turn on the power. At startup, the Infoblox logo appears in the LCD on the front panel of the appliance. Then the LCD scrolls repeatedly through a series of display screens.
2. To change the network settings for the LAN1 port, press one of the navigation buttons. The LCD immediately goes into the input mode, in which you can enter the IP address, netmask, and gateway for the LAN1 port.
3. Use the navigation buttons to enter an IP address, netmask, and gateway address for the LAN1 port.
4. Cable the LAN1 port of the NIOS appliance to a network as described in the installation guide that shipped with your product.

**Method 2 – Using the CLI**

You can use the Infoblox CLI to make an initial network configuration through the set network command. To access the CLI, make a direct serial connection from your management system.

1. Connect a console cable from the console port on your workstation to the male DB-9 console port on the NIOS appliance. The DB-9 pin assignments follow the EIA232 standard. You can use one RJ-45 rollover cable and two female RJ-45-to-female DB-9 adapters, or a female DB-9-to-female DB-9 null modem cable.

*Console Connection*



2. Use a serial terminal emulation program, such as Hilgraeve Hyperterminal® (provided with Windows® operating systems), to launch a session. The connection settings are as follows:
  - **Bits per second:** 9600
  - **Data bits:** 8

- **Parity:** None
  - **Stop bits:** 1
  - **Flow control:** Xon/Xoff
3. Log in to the appliance using the default username and password (**admin** and **infoblox**).
  4. At the Infoblox command prompt, enter **set network** to change the network settings, such as the IP address, netmask, and gateway for the LAN1 port. You can configure either IPv4 or IPv6 address for the LAN1 port. Note that In the following example, the variable `ip_addr1` is the IP address of the LAN1 port and `ip_addr2` is the IP address of the gateway for the subnet on which you set the `ip_addr1` address. Infoblox appliances support IPv4 and IPv6 networking configurations in all deployments cited in this chapter. You can set the LAN1 port to an IPv6 address and use that address to access the NIOS UI.

You can configure an IPv4 address for the LAN1 port as follows:

```
Infoblox > set network
```

```
NOTICE: All HA configuration is performed from the GUI. This interface is used only to configure a standalone node or to join a Grid.
```

```
Enter IP address: ip_addr1
```

```
Enter netmask: netmask
```

```
Enter gateway address: ip_addr2
```

```
Configure IPv6 network settings? (y or n): y
```

```
Enter IPv6 address [Default: none]: 2001:db8:a22:a00::29
```

```
Enter IPv6 Prefix Length [Default: none]: 64
```

```
Enter IPv6 gateway [Default: none] 2001:db8:a22:a00::1
```

```
  Become Grid member? (y or n): n
```

To avoid configuring IPv6 network settings, you can press **N** at the command line. You can configure an IPv6 address for the LAN1 port as follows:

```
Infoblox > set network
```

```
NOTICE: All HA configuration is performed from the GUI. This interface is used only to configure a standalone node or to join a grid.
```

```
Enter IP address : 2620:010A:6000:2400:0000:0000:0000:6508
```

```
Enter IPv6 Prefix Length: 64
```

```
Enter IPv6 gateway [Default: none]: 2620:010A:6000:2400:0000:0000:0000:0001
```

```
Configure IPv4 network settings? (y or n):n
```

```
Become grid member? (y or n): n
```

To configure IPv4 network settings, you can press **Y** at the command line and configure IPv4 address, netmask, and the gateway address. After you confirm your network settings, the Infoblox application automatically restarts. You can press **N** to avoid configuring IPv6 on the command line. After you confirm your network settings, the Infoblox application automatically restarts.

5. Cable the LAN1 port to a network. For information about installing and cabling the appliance, refer to the user guide or installation guide that was shipped with the product.

### Method 3 – Using the Infoblox NIOS Startup Wizard

When you first make an HTTPS connection to a NIOS appliance, the *Infoblox NIOS Startup Wizard* guides you through the deployment options and basic network settings. You can also change the password of the superuser (admin) and set up the system clock.

Note that you can configure network settings using the Startup wizard any time after you have configured the appliance. To start the wizard, from Grid Manager, select the **System** tab, and then click **System Properties** -> **Startup Wizard** from the Toolbar.

To make an HTTPS session to the appliance, you must be able to reach its IP address from the management system.



#### Note

If you have already set the IP address of the LAN1 port through the LCD or CLI so that you can reach it over the network—and you have already cabled the appliance to the network—you can skip the first step.

1. If you have not changed the default IP address (192.168.1.2/24) of the LAN1 port through the LCD or CLI—and the subnet to which you connect the appliance is not 192.168.1.0/24—put your management system in the 192.168.1.0/24 subnet and connect an Ethernet cable between the management system and the appliance.
2. Open an Internet browser window and enter **https://<IP\_address\_of\_the\_appliance>** to make an HTTPS connection. For information about supported browsers, see [Supported Browsers](#).  
Several certificate warnings may appear during the login process, because the preloaded certificate is self-signed and has the host name [www.infoblox.com](http://www.infoblox.com), which may not match the destination IP address you entered in step 1. To stop the warning messages from occurring each time you log in to Grid Manager, you can generate a new self-signed certificate or import a third-party certificate with a common name that matches the FQDN (fully qualified domain name) of the appliance. For information, see [Managing Certificates](#).
3. Enter the default username and password (**admin** and **infoblox**) on the Grid Manager login page, and then click **Login** or press **ENTER**. For information, see [Logging on to the NIOS UI](#).
4. Read the Infoblox End-User License Agreement (EULA), and then click **I Accept**.  
Note that if you want to view the privacy policy of Infoblox, then on the EULA screen, click **Infoblox Privacy Policy**. Grid Manager displays the policy on a new browser tab.
5. Read about the *Infoblox Customer Experience Improvement Program* and choose whether to participate (opt in) or not participate (opt out) in the program. By default, participation is enabled. If you want to opt out of the program, select **To Opt-Out of the alert program, please click here**. For more information about the program, see [Configuring the Customer Experience Improvement Program](#).
6. Click **OK**. The *Grid Setup* wizard appears.
7. In the first screen of the *NIOS Setup* wizard, complete the following:
8. **Type of Network Connectivity:** Select the type of network connectivity for the appliance from the drop-down list:
  - **IPv4 and IPv6:** Select this to configure a dual mode appliance.
  - **IPv4:** Select this to configure an IPv4 appliance.
  - **IPv6:** Select this to configure an IPv6 appliance.
9. **Are you configuring an HA pair or a standalone appliance?:** Select **Configuring a standalone appliance**. To configure an independent HA pair, see [Deploying an Independent HA Pair](#).

10. Click **Next** and complete the following to configure network settings:
  - **Host Name:** Enter a valid domain name for the appliance.
  - **Ports and Addresses:** This table lists the network interfaces based on the type of network connectivity of the appliance. For an IPv4 appliance, specify the network information for LAN1 (IPv4) interface and for an IPv6 appliance, specify the network information for LAN1 (IPv6) interface. For a dual mode appliance, specify the network information for both LAN1 (IPv4) and LAN1 (IPv6) interfaces.  
Enter correct information for the following by clicking the field:
    - **Interface:** Displays the name of the interface. You cannot modify this.
    - **IP Address:** Type the IPv4 or IPv6 address depending on the type of interface. An IPv6 address is a 128-bit number in colon hexadecimal notation. It consists of eight 16-bit groups of hexadecimal digits separated by colons (example: 2001:db8:0000:0123:4567:89ab:0000:cdef or 2001:db8::123:4567:89ab:0:cdef).
    - **Subnet Mask (IPv4) or Prefix Length (IPv6):** Specify an appropriate subnet mask for IPv4 address or prefix length for IPv6 address. The prefix length ranges from 2 to 127.
    - **Gateway:** Type the IPv4 or IPv6 address of the default gateway depending on the type of interface. For the IPv6 interface, you can also type **Automatic** to enable the appliance to acquire the IPv6 address of the default gateway and the link MTU from router advertisements.
    - **Port Settings:** Select the port settings from the drop-down list. The list contains all the settings supported by the hardware model. The default is **Automatic**. The appliance automatically detects the port settings.
11. Click **Next** and complete the following to set admin password:
  - **Yes:** To change the default password.
  - **No:** To keep the default password. Infoblox recommends that you change the default password. When you select **Yes**, complete the following:
    - **Password:** Enter a password for the superuser admin account. The password must be a single alphanumeric string without spaces and at least four characters long. The password is case-sensitive.
    - **Retype Password:** Enter the same password.
12. Click **Next** and complete the following to configure time settings:
  - **Time Zone:** Select the applicable time zone from the drop-down list. The default is **(UTC) Coordinated Universal Time**.
  - **Would you like to enable NTP?:**
    - Select **Yes** to synchronize the time with external NTP servers, and then click the Add icon. Grid Manager adds a row to the NTP Server table. Click the row and enter either the IP address (IPv4 or IPv6) or the resolvable host name of an NTP server. You can view a list of public NTP servers at ntp.isc.org.
    - Select **No** to specify the time settings for the appliance.
  - **Date:** Enter the date in YYYY-MM-DD format. You can also click the calendar icon to select a date from the calendar widget.
  - **Time:** Enter the time in HH:MM:SS AM/PM format. You can also click the clock icon to select a time from the drop-down list.
13. Click **Next** to view the summary of the configuration. Review the information and verify that it is correct. You can change the information that you entered by clicking **Previous** to go back to a previous step.
14. Click **Finish**. The appliance restarts and disconnects Grid Manager.

In a dual mode independent appliance, the communication protocol for all the services is set to IPv4 by default. You can change the default communication protocol for the services.

### Changing the Communication Protocol for a Dual Mode Independent Appliance

You can change the default communication protocol for a dual mode appliance. You can force the appliance to use a specific protocol for the reporting services and for services with two types of resolution (A and AAAA records), you can set the preferred communication protocol.

To change the communication protocol for a dual mode independent appliance, complete the following:

1. From the **System** tab, select the **System Manager** tab -> click **System Properties** from the Toolbar, and select **Edit** from the drop-down list.
2. In the *System Properties* editor, select the **Network** tab -> **Basic** tab, and then complete the following:



- **Communication Protocol Settings and Preferences:** Select either **IPv4** or **IPv6** from the drop-down list. This setting will force the appliance to use the specified protocol for the reporting services and this is the preferred protocol for services with two types of resolution (A and AAAA records).
- **Customized Settings:** Select this and perform the following:
  - **Always use this Communications Protocol for:** Select either **IPv4** or **IPv6** from the drop-down list. This setting will force the appliance to use the specified communication protocol for reporting services.
  - **Always prefer this Communications Protocol for:** Select either **IPv4** or **IPv6** from the drop-down list to specify the preferred communication protocol for the listed services, which has two types of resolution (A and AAAA records). The appliance uses the preferred protocol first for the service.

## Configuration Example: Deploying a NIOS Appliance as a Primary DNS Server

In this example, you will configure the NIOS appliance as a primary DNS server for corpxyz.com. Its FQDN (fully-qualified domain name) is ns1.corpxyz.com. The interface IP address of the LAN1 port is 10.1.5.2/24. Because this is a private IP address, you must also configure the firewall to perform NAT (network address translation), mapping the public IP address 1.1.1.2 to 10.1.5.2. Using its public IP address, ns1 can communicate with appliances on the public network.

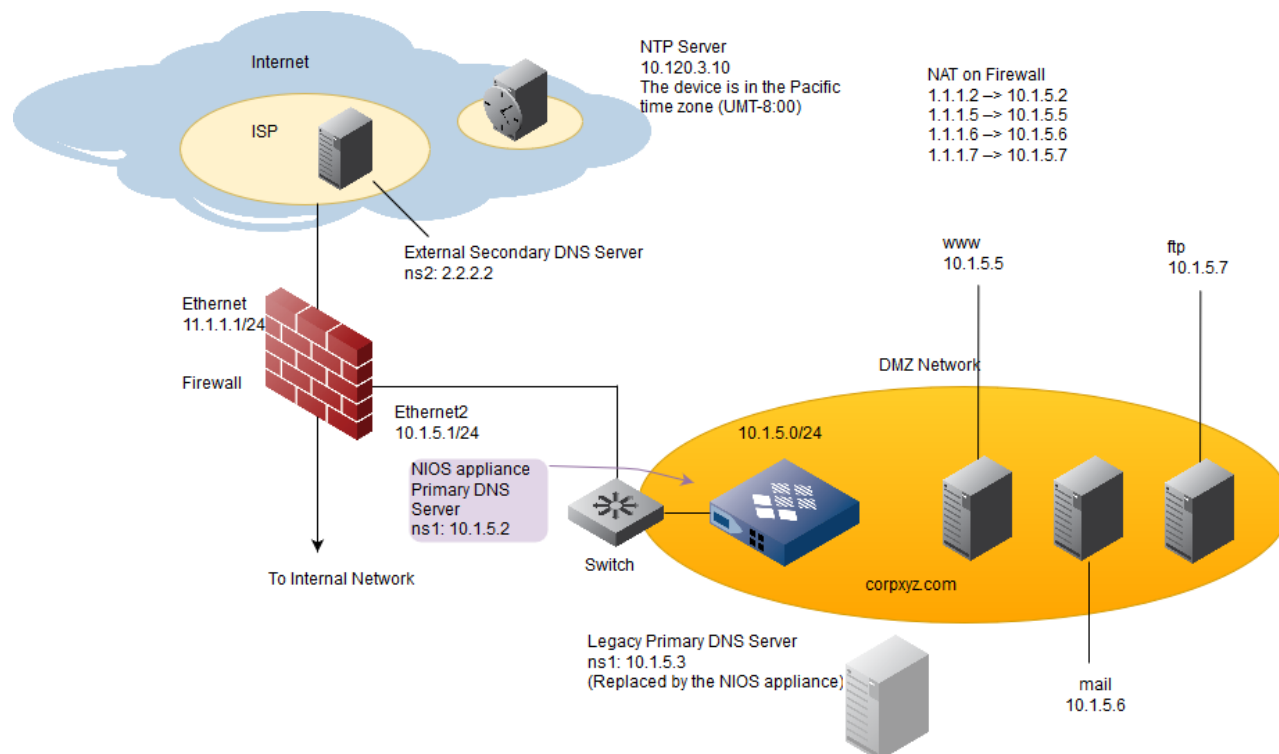
The FQDN and IP address of the external secondary DNS server are ns2.corpxyz.com and 2.2.2.2. The ISP hosts this server. You can deploy NIOS appliance in IPv4, IPv6, or dual mode (IPv4 and IPv6), but the configuration example uses IPv4 addresses.

The primary and secondary servers answer queries for the following public-facing servers in the DMZ:

- www.corpxyz.com
- mail.corpxyz.com
- ftp.corpxyz.com

When you create the corpxyz.com zone on the NIOS appliance, you import zone data from the legacy DNS server at 10.1.5.3.

*Example 1 Network Diagram*



## Cabling the Appliance to the Network and Turning On Power

Connect an Ethernet cable from the LAN1 port of the NIOS appliance to a switch in the DMZ network and turn on the power. For information about installing and cabling the appliance, refer to the user guide or installation guide that ships with the product.

## Specifying Initial Network Settings

Before you can configure the NIOS appliance through Grid Manager, you must be able to make a network connection to it. The default network settings of the LAN1 port are 192.168.1.2/24 with a gateway at 192.168.1.1 (the HA and MGMT ports do not have default network settings). To change these settings to suit your network, use either the LCD or the console port.

In this example, you change the IP address/netmask of the LAN1 port to 10.1.5.2/24, and the gateway to 10.1.5.1.

### LCD

The NIOS appliance has an LCD and navigation buttons on its front panel.

At startup, the Infoblox logo appears in the LCD on the front panel of the appliance. Then the LCD scrolls repeatedly through a series of display screens.

1. To change the network settings from the default, press one of the navigation buttons.  
The LCD immediately goes into input mode, in which you can enter the IP address, netmask, and gateway for the LAN1 port.
2. Use the navigation buttons to enter the following information:
  - IP Address: 10.1.5.2
  - Netmask: 255.255.255.0
  - Gateway: 10.1.5.1

## Specifying Appliance Settings

When you make the initial HTTPS connection to the NIOS appliance, the NIOS Startup Wizard guides you through the basic deployment of the appliance on your network. Use the wizard to enter the following information:

- Deployment: single independent appliance
- Host name: ns1.corpxyz.com
- Password: SnD34n534
- NTP (Network Time Protocol) server: 10.120.3.10; time zone: (UMT – 8:00 Pacific Time (US and Canada), Tijuana

1. Open an Internet browser window and enter <https://10.1.5.2>.
2. Accept the certificate when prompted.  
Several certificate warnings may appear during the login process. This is normal because the preloaded certificate is self-signed and has the hostname [www.infoblox.com](https://www.infoblox.com), which does not match the destination IP address you entered in step 1. To stop the warning messages from occurring each time you log in to Grid Manager, you can generate a new self-signed certificate or import a third-party certificate with a common name that matches the FQDN (fully-qualified domain name) of the appliance. This is a very simple process. For information about certificates, see [Creating a Login Banner](#).
3. Enter the default username and password (**admin** and **infoblox**) on the Grid Manager login page, and then click **Login** or press **ENTER**. For information, see [Logging on to the NIOS UI](#).
4. Read the Infoblox End-User License Agreement, and then click **I Accept** to proceed.
5. Read about the *Infoblox Customer Experience Improvement Program* and choose whether to participate (opt in) or not participate (opt out) in the program. By default, participation is enabled. If you want to opt out of the program, select **To Opt-Out of the alert program, please click here**. For more information about the program, see [Configuring the Customer Experience Improvement Program](#).

6. Click **OK**. Grid Manager may take a few seconds to load your user profile.
7. In the first screen of the *NIOS Startup* wizard, complete the following:
  - **Type of Network Connectivity**: Select **IPv4** from the drop-down list.
  - **Are you configuring an HA pair or a standalone appliance?**: Select **Configuring a standalone appliance**. To configure an independent HA pair, see [Deploying an Independent HA Pair](#).
8. Click **Next** and complete the following to configure network settings:
  - **Host Name**: Enter **ns1.corpxyz.com**.
  - **Ports and Addresses**: Specify the network settings for LAN1 (IPv4) port. Enter correct information for the following by clicking the field:
    - **IP Address**: Enter **10.1.5.2** as the IPv4 address for the LAN1 port.
    - **Subnet Mask (IPv4) or Prefix Length (IPv6)**: Enter **255.255.255.0** as the subnet mask for the LAN1 (IPv4) port.
    - **Gateway**: Enter **10.1.5.1** as the gateway of the subnet on which the LAN1 port is set.
    - **Port Settings**: Use the default value **Automatic**.
9. Click **Next** and complete the following to set admin password:
  - **Would you like to set admin password?**: Click **Yes**.
  - **Password**: Enter **SnD34n534**.
  - **Retype Password**: Enter **SnD34n534** again.
10. Click **Next** and complete the following to configure the time settings:
  - **Time Zone**: Select **UMT – 8:00 Pacific Time (US and Canada), Tijuana** from the drop-down list.
  - **Would you like to enable NTP?**: Select **Yes** to synchronize the time with external NTP servers, and then click the Add icon. Grid Manager adds a row to the NTP Server table. Click the row and enter **10.120.3.10** in the **NTP Server** field.
11. Click **Next** to view the summary of the configuration. Review the information and verify that it is correct. You can change the information you entered by clicking **Previous** to go back to a previous step.
12. Click **Finish**.

## Enabling Zone Transfers on the Legacy Name Server

To allow the appliance to import zone data from the legacy server 10.1.5.3, you must configure the legacy server to allow zone transfers to the appliance at 10.1.5.2.

### Legacy BIND Server

1. Open the `named.conf` file using a text editor and change the `allow-transfer` statement as shown below:
  - **For All Zones** — To set the `allow-transfer` statement as a global statement in the `named.conf` file for all zones:
 

```
options {
zone-statistics yes;
directory "/var/named/named_conf"; version "";
recursion yes;
listen-on { 127.0.0.1; 10.1.5.3; };
...
allow-transfer {10.1.5.2; }; transfer-format many-answers;
};
```
  - **For a Single Zone** — To set the `allow-transfer` statement in the `named.conf` file for the `corpxyz.com` zone:
 

```
zone "corpxyz.com" in { type master;
allow-transfer {10.1.5.2;};
notify yes;
};
```
2. After editing the `named.conf` file, restart DNS service on the appliance for the change to take effect.

## Legacy Windows 2000/2003 Server

1. Click **Start** -> **All Programs** -> **Administrative Tools** -> **DNS**.
2. Click **+** (for ns1) -> **+** (for Forward Lookup Zones) -> **corpxyz.com**.
3. Right-click **corpxyz.com**, and then select **Properties** -> **Zone Transfers**.
4. On the *Zone Transfers* page in the *corpxyz.com Properties* dialog box, enter the following:
  - **Allow zone transfers**: Select this.
  - **Only to the following servers**: Select this.
  - **IP address**: Enter **10.1.5.2**, and then click **Add**.
5. To save the configuration and close the *corpxyz.com Properties* dialog box, click **OK**.

## Importing Zone Data on an Independent Appliance

You can import zone data from a legacy server or manually enter it. When you import both forward-mapping and reverse-mapping zone data, the NIOS appliance automatically creates Infoblox host records if corresponding A and PTR records are present. You can then modify the host records to add MAC addresses. However, if you only import forward-mapping zone data, the NIOS appliance cannot create host records from just the A records. In that case, because you cannot later convert A records to host records, it is more efficient to create the corpxyz.com zone, and define host records manually.

Infoblox host records are data models that represent IP devices within the Infoblox semantic database. The NIOS appliance uses a host object to define A, PTR, and CNAME resource records in a single object, as well as a DHCP fixed address if you include a MAC address in the host object definition. The host object prevents costly errors because you only maintain a single object for multiple DNS records and a DHCP fixed address. Therefore, it is advantageous to use host records instead of separate A, PTR, and CNAME records.



### Note

If you only have forward-mapping zones on your legacy servers and you want to add reverse-mapping zones, automatically convert A records to host records in the imported forward-mapping zones, and create reverse host records in corresponding reverse-mapping zones, create the reverse-mapping zones on the NIOS appliance and then import the forward-mapping zones data. The NIOS appliance automatically converts the imported A records to host records in the forward-mapping zones and creates reverse host records in the reverse-mapping zones.

You also have the option of using the *Data Import Wizard* for loading DNS and DHCP data. For large data sets, this option is an efficient approach. To download the *Data Import Wizard*, visit [www.infoblox.com/import/](http://www.infoblox.com/import/).

In this example, when you create the corpxyz.com forward-mapping zone, you import zone data for the existing corpxyz.com zone from the legacy server at 10.1.5.3. When you create the 1.1.1.0/24 reverse-mapping zone, you also import the reverse-mapping zone records from the legacy server. After the appliance has both the forward- and reverse-mapping zone data, it converts the A and PTR records to Infoblox host records.

## Creating a Name Server Group

1. Open an Internet browser window, enter <https://10.1.5.2>, and then log in to Grid Manager using the username **admin** and password **SnD34n534**.
2. From the **Data Management** tab, select the **DNS** tab -> **Name Server Groups** tab, and then click the Add icon -> **Name Server Group**.
3. In the *Name Server Group* wizard, complete the following:
  - **Name**: Enter **corpxyz** as the group name.
  - **Name Servers**: Click the Add icon -> **Primary**.
  - In the *Add Primary* section, Grid Manager displays the host name of the independent appliance. Click **Add**.  
Grid Manager adds the independent system as the primary server.
  - Click the Add icon -> **External Secondary**.

- In the *Add External Secondary* section, complete the following:
    - **Name:** Enter **ns2.corpxyz.com**.
    - **Address:** Enter **2.2.2.2**.
    - **Stealth:** Clear this checkbox.
    - Click **Add**. Grid Manager adds the external secondary to the name server group.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Creating a Forward-Mapping Zone



### Note

To import zone data, you must first create a zone and save it.

1. To create an authoritative zone, from the **Data Management** tab, select the **DNS** tab -> **Zones** tab, and then click the Add icon -> **Authoritative Zone**.
2. In the *Add Authoritative Zone* wizard, select **Add an authoritative forward-mapping zone**.
3. Click **Next** and complete the following:
  - **Name:** Enter **corpxyz.com**.
  - **Comment:** Enter **DNS zone**.
4. Click **Next** to assign a name server group to the zone.
5. Click the **Zones** tab, select the **corpxyz.com** checkbox, and then click the Edit icon.
6. In the *Authoritative Zone* editor, select the **Name Servers** tab, and then complete the following:
  - **Use this name server group:** Select this, and then select **corpxyz** from the drop-down list.
7. Save the configuration and click **Restart** if it appears at the top of the screen.

## Importing Zone Data

1. To import zone data to the corpxyz.com zone that you created earlier, click the **Zones** tab, select the **corpxyz.com** checkbox, and then click **Import Zone** from the Toolbar.
2. In the *Import Zone* editor, complete the following:
  - **Address:** Enter the IP address **10.1.5.3** of the DNS server from which you want to import zone data.
  - **Create Hosts and Bulk Hosts during import:** Select this option to allow the appliance to merge imported records into hosts and bulk hosts. If you do not select this option, then resource records are imported one-to-one with each DNS record from the imported zone producing a corresponding DNS record on the NIOS appliance.
    - **Create PTR records for Hosts if necessary:** Select this to create host records from the imported address records, even if the corresponding PTR records or the authoritative reverse zones that would contain them, do not exist. If you do not select this option, then host records will be created for imported address records only if a corresponding PTR record exists.
    - **Create PTR records for Bulk Hosts if necessary:** Select this to create bulk host records from the imported address records, even if the corresponding PTR records or the authoritative reverse zones that would contain them, do not exist. If you do not select this option, then bulk host records will be created for the imported address records only if the corresponding PTR records exist.
3. Click **Import**.
4. After successfully importing the zone data, click **corpxyz.com** in the **Zones** tab.  
You can see all the imported forward-mapping zone data in the *Records* panel. Because you have not yet imported the reverse-mapping zone data, most of the records appear as A records.
5. To import the reverse-mapping zone data, from the **Zones** tab, click the Add icon -> **Authoritative Zone**.
6. In the *Add Authoritative Zone* wizard, select **Add an authoritative IPv4 reverse-mapping zone**.
7. Click **Next** and complete the following:
  - **IPv4 Network:** Enter **1.1.1.0**.
  - **Netmask:** Select **24** from the drop-down list.
  - **Comment:** Enter **Reverse-mapping zone**.
8. Click **Save & Close**.
9. To assign a name server group to the reverse-mapping zone, click the **Zones** tab, select the **1.1.1.in-addr.arpa** checkbox, and then click the Edit icon.
10. In the *Authoritative Zone* editor, select the **Name Servers** tab, and then complete the following:
  - **Use this name server group:** Select this, and then select **corpxyz** from the drop-down list.

11. Click **Save & Close**.
12. To import reverse-mapping zone data, click the **Zones** tab, select the `corpxyz.com` checkbox, and then click **Import Zone** from the Toolbar.
13. In the *Import Zone* editor, complete the following:
  - **Address:** Enter the IP address `10.1.5.3` of the DNS server from which you want to import zone data.
14. Click **Import**.
15. After successfully importing the zone data, click `1.1.1.in-addr.arpa` on the **Zones** tab. You can see all the imported reverse-mapping zone data in the *Records* panel.
16. Click `corpxyz.com` in the Forward Mapping Zones list. Because you have now imported both the forward- and reverse-mapping zone data, most of the records appear as host records.
17. Finally, you must remove the `ns1` host record for the legacy server (value `1.1.1.3`). To remove it, select the `ns1` checkbox (the host record for `1.1.1.3`), and then click the Delete icon.

### Designating the New Primary on the Secondary Name Server (at the ISP Site)

In this example, the external secondary name server is maintained by an ISP, so you must contact your ISP administrator to change the IP address of the primary (or *master*) name server. (If you have administrative access to the secondary name server, you can make this change yourself.)

Because a firewall performing NAT exists between the secondary and primary name servers, specify the NAT address `1.1.1.2` for the primary name server instead of `10.1.5.2`.

### Secondary BIND Server

1. Open the `named.conf` file using a text editor and set `ns1` (with NAT address `1.1.1.2`) as the primary from which `ns2` receives zone transfers in the `named.conf` file for the `corpxyz.com` zone.
2. After editing the `named.conf` file, restart DNS service for the change to take effect.

### Secondary Windows 2000/2003 Server

1. Click **Start** -> **All Programs** -> **Administrative Tools** -> **DNS**.
2. Click **+** (for `ns2`) -> **+** (for Forward Lookup Zones) -> `corpxyz.com`.
3. Right-click `corpxyz.com`, and then select **Properties** -> **General**.
4. On the *General* page in the *corpxyz.com Properties* dialog box, enter the following:
  - **Zone file name:** `corpxyz.com.dns`
  - **IP address:** Enter `1.1.1.2` and then click **Add**.
  - In the IP Address field, select `1.1.1.3` (the NAT IP address of the legacy DNS server), and then click **Remove**.
5. To save the configuration and close the *corpxyz.com Properties* dialog box, click **OK**.

### Configuring NAT and Policies on the Firewall

Change the NAT and policy settings on the firewall to allow bidirectional DNS traffic to and from `ns1.corpxyz.com` and NTP traffic from `ns1.corpxyz.com` to the NTP server at `10.120.3.10`.

For example, enter the following commands on a Juniper firewall running ScreenOS 4.x or later:

```
set address dmz ns1 10.1.5.2/32
set address untrust ntp_server 10.120.3.10/32 set interface ethernet1 mip
1.1.1.2 host 10.1.5.2 set policy from dmz to untrust ns1 any dns permit
set policy from untrust to dmz any mip(1.1.1.2) dns permit set policy from dmz
to untrust ns1 ntp_server ntp permit
```

At this point, the new DNS server can take over DNS service from the legacy server. You can remove the legacy server and unset any firewall policies permitting traffic to and from 10.1.5.3.

## Deploying an Independent HA Pair

To deploy an independent HA pair, you cable the HA and LAN1, LAN1 (VLAN), or LAN2, LAN2 (VLAN) ports to the network and configure the IP settings for these ports and the VIP address within the same subnet. For more information about HA pairs, see [About HA Pairs](#).

The default LAN1 or LAN2 settings are as follows:

- IP address: 192.168.1.2
- Netmask: 255.255.255.0
- Gateway: 192.168.1.1

You can configure an HA pair using the *Infoblox NIOS Startup Wizard*. IPv4 and IPv6 network values are supported for the *NIOS Startup Wizard* and for HA Pair configuration. The NIOS appliance MGMT port also can be configured to support an IPv6 address.

- Requirements: HTTPS connections from your management system to the Ethernet ports on the two appliances.
- Advantage: The startup wizard provides step-by-step guidance for configuring the network settings of the VIP address and HA and LAN1 or LAN1 (VLAN) ports on both nodes, for setting the host name, admin password, and system clock, and—if using NTP (Network Time Protocol)—for enabling the HA pair as an NTP client.

### Using the Infoblox NIOS Startup Wizard to Configure an HA Pair

When you first make an HTTPS connection to the NIOS appliance, the Infoblox NIOS Startup Wizard guides you through various deployment options, basic network settings, and opportunities for changing the password of the superuser *admin* and for setting the system clock.

### Configuring the Connecting Switch

To ensure that VRRP (Virtual Router Redundancy Protocol) works properly, configure the following settings at the port level for all the connecting switch ports (HA, LAN1, LAN1 (VLAN), LAN2, and LAN2 (VLAN)):

- Spanning Tree Protocol: Disable. For vendor specific information, search for "HA" in the Infoblox Knowledge Base system at <https://support.infoblox.com>.
- Trunking: Disable
- EtherChannel: Disable
- IGMP Snooping: Disable
- Port Channeling: Disable
- Speed and Duplex settings: Match these settings on both the Infoblox appliance and switch
- Disable other dynamic and proprietary protocols that might interrupt the forwarding of packets



#### Note

By default, a NIOS appliance automatically negotiates the optimal connection speed and transmission type (full or half duplex) on the physical links between its LAN1, HA, and MGMT ports and the Ethernet ports on the connecting switch. If the two appliances fail to auto-negotiate the optimal settings, see [Modifying Ethernet Port Settings](#) for steps you can take to resolve the problem.

### Putting Both Nodes on the Network

1. Use one of the methods described in [Deploying a Single Independent Appliance](#) to configure the network settings of the LAN1 port of each node so that they are on the same subnet and you can reach them across the network.



2. Cable the LAN1 port and the HA port on each node to the network switch.
3. Cable your management system to the network switch.

## Configuring Node 1

1. Open an Internet browser window and enter `https://* <the IP address of the appliance>` to make an HTTPS connection to the first node. For information about supported browsers, see [Supported Browsers](#). Several certificate warnings may appear during the login process because the preloaded certificate is self-signed and has the hostname `www.infoblox.com`, which may not match the destination IP address that you entered in step 1. To stop the warning messages from occurring each time you log in to Grid Manager, you can generate a new self-signed certificate or import a third-party certificate with a common name that matches the FQDN (fully qualified domain name) of the appliance. For information, see [Creating a Login Banner](#).
2. Enter the default username and password (**admin** and **infoblox**) on the Grid Manager login page, and then click **Login** or press **ENTER**. For information, see [Logging on to the NIOS UI](#).
3. Read the Infoblox End-User License Agreement, and then click **I Accept** to proceed.
4. Read about the *Infoblox Customer Experience Improvement Program* and choose whether to participate (opt in) or not participate (opt out) in the program. By default, participation is enabled. If you want to opt out of the program, select **To Opt-Out of the alert program, please click here**. For more information about the program, see [Configuring the Customer Experience Improvement Program](#).
5. Click **OK**. Grid Manager may take a few seconds to load your user profile.
6. In the first screen of the *Grid Setup* wizard, select **Configure a Grid Master** and click **Next**.
7. Specify the following:
  - **Grid Name:** Type the name of the Grid. The default name is *Infoblox*.
  - **Shared Secret:** Enter the shared secret that both nodes use to authenticate each other when establishing a VPN tunnel for ensuing bloxSYNC traffic. The default shared secret is **test**.
  - **Confirm Shared Secret:** Enter the shared secret again.
  - **Host Name:** Enter a valid domain name for the node.
  - **Type of Network Connectivity:** Select the type of network connectivity from the drop-down list:
    - **IPv4 and IPv6:** Select this to configure a dual mode HA pair.
    - **IPv4:** Select this to configure an IPv4 HA pair.
    - **IPv6:** Select this to configure an IPv6 HA pair.
  - Select **Yes** in the **Is the Grid Master an HA pair** field for the first appliance of the HA pair.
    - **Send HA and Grid Communication over:** Select either **IPv4** or **IPv6**. This field is displayed only when you configure a dual mode (IPv4 and IPv6) HA pair.
8. Click **Next** and complete the following to set properties for the first node:
  - **Virtual Router ID:** Enter the VRID (virtual router ID). This must be a unique VRID number—from 1 to 255—for this subnet.
  - **Ports and Addresses:** This table lists the network interfaces depending on the type of network connectivity. For IPv4 HA pair, specify the network information for VIP (IPv4), Node1 HA (IPv4), Node2 HA (IPv4), Node1 LAN1 (IPv4), and Node2 LAN1 (IPv4) interfaces. For IPv6 HA pair, specify the network information for VIP (IPv6), Node1 LAN1 (IPv6), and Node2 LAN1 (IPv6) interfaces. For a dual mode HA pair, if you select **IPv4** in the **Send HA and Grid Communication over** field in step 2 of the *NIOS Startup* wizard, specify the network information for the following interfaces: VIP (IPv4), Node1 HA (IPv4), Node1 LAN1 (IPv4), Node2 HA (IPv4), Node2 LAN1 (IPv4), VIP (IPv6), Node1 LAN1 (IPv6), and Node2 LAN1 (IPv6) ports. If you select **IPv6** in the **Send HA and Grid Communication over** field in step 2 of the *NIOS Startup* wizard, specify the network information for the following interfaces: VIP (IPv4), Node1 LAN1 (IPv4), Node2 LAN1 (IPv4), VIP (IPv6), Node1 LAN1 (IPv6), and Node2 LAN1 (IPv6). Click the empty fields and complete the following information:
    - **Interface:** Displays the name of the interface. You cannot modify this.
    - **Address:** Type the IPv4 or IPv6 address depending on the type of interface.
    - **Subnet Mask (IPv4) or Prefix Length (IPv6):** Specify an appropriate subnet mask for IPv4 address or prefix length for IPv6 address. The prefix length ranges from 2 to 127.
    - **Gateway:** Type the IPv4 or IPv6 address of the default gateway depending on the type of interface. For IPv6 interface, you can also type **Automatic** to enable the appliance to acquire the IPv6 address of the default gateway and the link MTU from router advertisements.
    - **Port Settings:** Select the port settings from the drop-down list. The list contains all settings supported by the hardware model. The default is **Automatic**. The appliance automatically detects the port settings.



9. Click **Next** and complete the following to set admin password:
  - **Yes:** To change the default password.
  - **No:** To keep the default password.
 If you select **Yes**, complete the following:
  - **Password:** Enter a password for the superuser admin account. The password cannot contain spaces and it must be at least four characters long. The password is case-sensitive.
  - **Retype Password:** Enter the same password.
10. Click **Next** and complete the following to configure time settings:
  - **Time Zone:** Select the applicable time zone from the drop-down list. The default is **(UTC) Coordinated Universal Time**.
  - **Would you like to enable NTP?:**
    - Select **Yes** to synchronize the time with external NTP servers. Click the Add icon. Grid Manager adds a row to the NTP Server table. Click the row and enter either the IPv4 or IPv6 address or the resolvable host name of an NTP server. You can view a list of public NTP servers at [ntp.isc.org](http://ntp.isc.org).
    - Select **No** to specify a date and time.
      - **Date:** Enter the data in YYYY-MM-DD format. You can also click the calendar icon to select a date from the calendar widget.
      - **Time:** Enter the time in HH:MM:SS AM/PM format. You can also click the clock icon to select a time from the drop-down list.
11. Click **Next** to view the summary of the configuration. Review the information and verify that it is correct. You can change the information you entered by clicking **Previous** to go back to a previous step.
12. Click **Finish**.

## Configuring Node 2

1. Open an Internet browser window and enter **https://\* <the IP address of the appliance>** to make an HTTPS connection to the second node. For information about supported browsers, see [Supported Browsers](#). Several certificate warnings may appear during the login process because the preloaded certificate is self-signed and has the host name [www.infoblox.com](http://www.infoblox.com), which may not match the destination IP address you entered in step 1. To stop the warning messages from occurring each time you log in to Grid Manager, you can generate a new self-signed certificate or import a third-party certificate with a common name that matches the FQDN (fully qualified domain name) of the appliance. For more information, see [Creating a Login Banner](#).
2. Enter the default username and password (**admin** and **infoblox**) on the Grid Manager login screen, and then click **Login** or press **ENTER**. For more information, see [Logging on to the NIOS UI](#).
3. Read the Infoblox End-User License Agreement, and then click **I Accept** to proceed.
4. Read about the [Infoblox Customer Experience Improvement Program](#) and choose whether to participate (opt in) or not participate (opt out) in the program. By default, participation is enabled. If you want to opt out of the program, select **To Opt-Out of the alert program, please click here**. For more information about the program, see [Configuring the Customer Experience Improvement Program](#).
5. Click **OK**. Grid Manager may take a few seconds to load your user profile.
6. In the first screen of the *NIOS Setup* wizard, complete the following:
  - **Type of Network Connectivity:** Select the type of network connectivity from the drop-down list:
    - **IPv4 and IPv6:** Select this to configure a dual mode HA pair.
    - **IPv4:** Select this to configure an IPv4 HA pair.
    - **IPv6:** Select this to configure an IPv6 HA pair.
  - Select **Configuring an HA pair** to configure an independent HA pair and click **No** to configure the first node of an HA pair
7. Click **Next** and complete the following to configure network settings:
  - **HA Virtual IP address:** Enter the VIP (virtual IP) address and its netmask.
  - **HA Pair Name:** Enter a name for the HA pair. The default name is **Infoblox**. Ensure that you use the same name as the first node.
  - **Shared Secret:** Enter a text string that both nodes use as a shared secret to authenticate each other when establishing a VPN tunnel. The default shared secret is a test. This must be the same shared secret that you entered on the first appliance.
  - **Show Password:** Click this to display the shared secret. Clear it to conceal the shared secret.
8. Click **Next**, and then complete the following to set properties for the second appliance:
  - **IP Address:** Enter the IPv4 or IPv6 address of the appliance.

- **Subnet Mask:** Enter the subnet mask of the appliance.  
Or
  - **Prefix Length:** Enter the prefix length if you have entered the IPv6 address in the **IP Address** field. The prefix length ranges from 2 to 127.
  - **Gateway:** Enter the IP address of the gateway of the subnet of the interface.
9. Click **Next** to view the summary of the configuration. Review the information and verify that it is correct. You can change the information you entered by clicking **Previous** to go back to a previous step.
  10. Click **Finish**.

The setup of the HA pair is complete. When you next make an HTTPS connection to the HA pair, use the VIP address.

The communication protocol for all the services in a dual mode (IPv4 and IPv6) HA appliance is the same protocol as the one used for VRRP advertisements. For example, if you select **IPv4** in the **Send HA and Grid Communication Over** field on the first screen of the *NIOS Setup* wizard, then IPv4 is set as the communication protocol for all the services. However, you can override the communication protocol for all the services in a dual mode HA pair. For information, see [Changing the Communication Protocol for a Dual Mode Independent Appliance](#).

## Configuration Example: Configuring an HA Pair for Internal DNS and DHCP Services

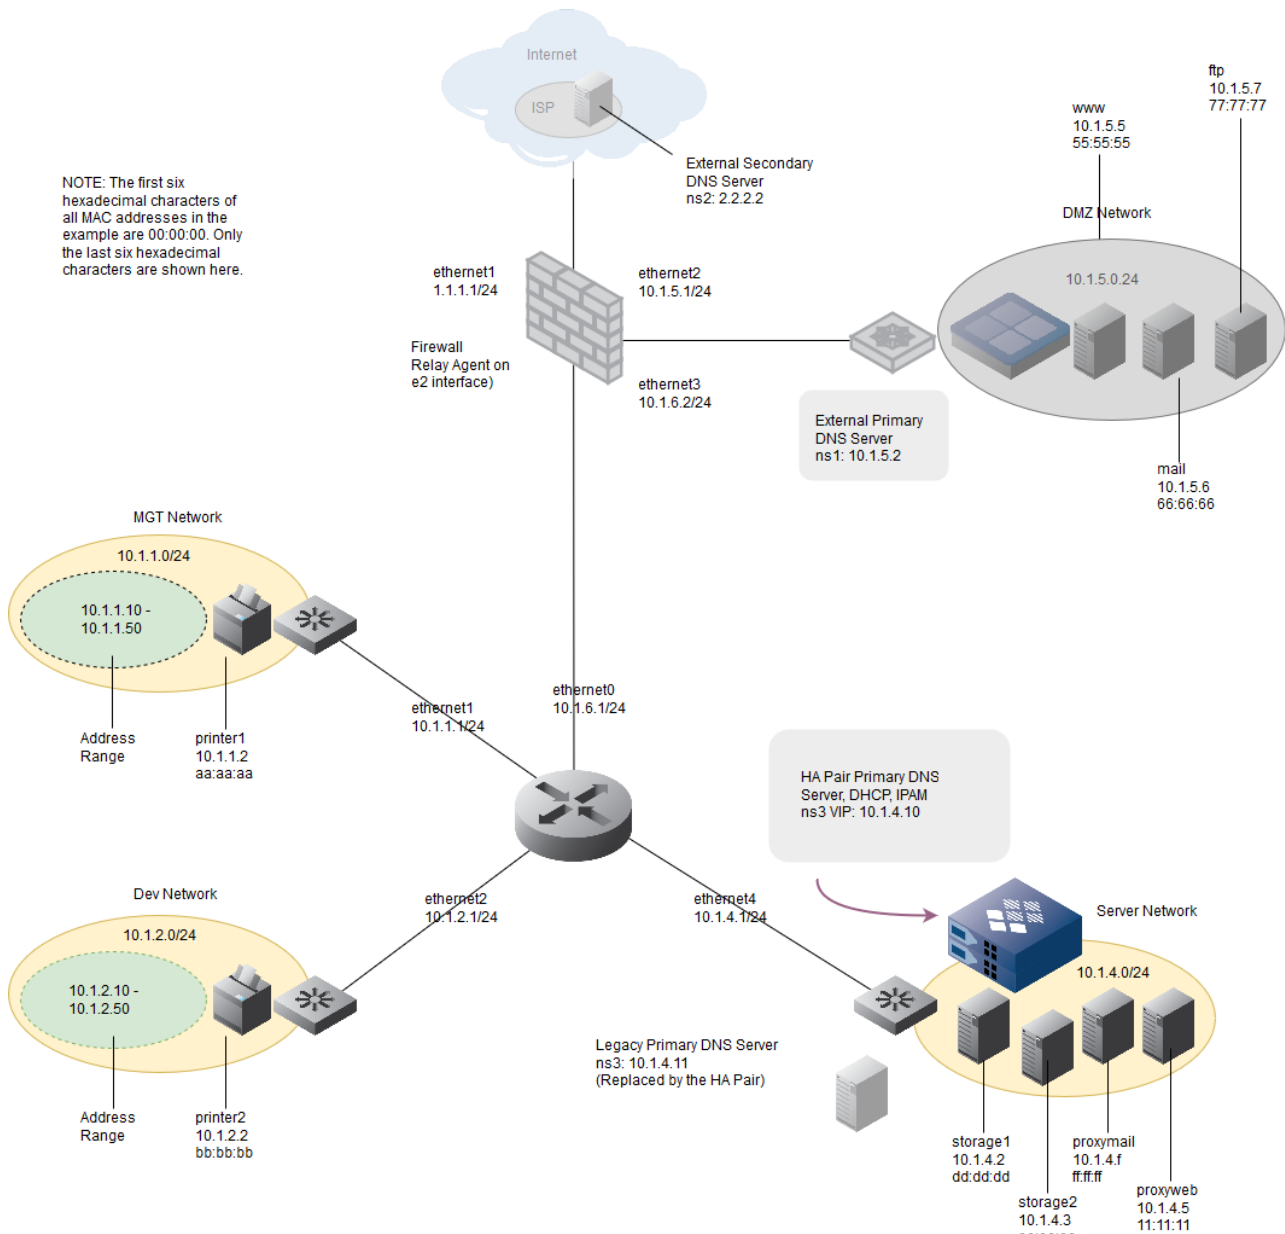
In this example, you set up an HA pair of NIOS appliances to provide internal DNS and DHCP services. The HA pair answers internal queries for all hosts in its domain (corpxyz.com). It forwards internal queries for external sites to ns1.corpxyz.com at 10.1.5.2 and ns2.corpxyz.com at 2.2.2.2. It also uses DHCP to provide dynamic and fixed addresses. You can deploy the HA appliance in IPv4, IPv6 or dual mode(IPv4 and IPv6), but the configuration example uses IPv4 addresses.

The HA pair consists of two appliances (nodes). The IP addresses of the VIP (virtual IP) address of the HA pair and the HA and LAN1 ports on each node are as follows:

HA Pair IP Addresses	
VIP 10.1.4.10 (the address that the active node of the HA pair uses)	
Node 1	Node 2
<ul style="list-style-type: none"> <li>• LAN1 10.1.4.6</li> <li>• HA 10.1.4.7</li> </ul>	<ul style="list-style-type: none"> <li>• LAN1 10.1.4.8</li> <li>• HA 10.1.4.9</li> </ul>

The virtual router ID number for the HA pair is 150. The ID number must be unique for this network segment. When you create the corpxyz.com zone on the HA pair, you import DNS data from the legacy server at 10.1.4.11.

*Example 2 Network Diagram*



An HA pair of NIOS appliances provide internal DNS services. It answers internal queries for all hosts in its domain. It forwards internal queries for external sites to ns1 and ns2. It also serves DHCP, providing both dynamic and fixed addresses.

### Cabling Appliances to the Network and Turning On Power

Connect Ethernet cables from the LAN1 and HA ports on both NIOS appliances to a switch in the server network and turn on the power for both appliances. For information about installing and cabling the appliance, refer to the user guide or installation guide that ships with the product.

### Specifying Initial Network Settings

Before you can configure the appliances through Grid Manager, you must be able to make a network connection to them. The default network settings of the LAN1 port are 192.168.1.2/24 with a gateway at 192.168.1.1 (the HA and MGMT

ports do not have default network settings). To change these settings, you can use the LCD or make a console connection to each appliance.

#### Node 1

Using the LCD or console port on one of the appliances, enter the following information:

- IP Address: **10.1.4.6** (for the LAN1 port)
- Netmask: **255.255.255.0**
- Gateway: **10.1.4.1**

#### Node 2

Using the LCD or console port on the other appliance, enter the following information:

- IP Address: **10.1.4.8** (for the LAN1 port)
- Netmask: **255.255.255.0**
- Gateway: **10.1.4.1**

After you confirm your network settings, the Infoblox GUI application automatically restarts.

### Specifying Appliance Settings

When you make the initial HTTPS connection to a NIOS appliance, the Infoblox NIOS Startup Wizard guides you through the basic deployment of the appliance on your network. To set up an HA pair, you must connect to and configure each appliance individually.

#### Node 1

1. Open an Internet browser window and enter <https://10.1.4.6>.
2. Accept the certificate when prompted. Several certificate warnings may appear during the login process. This is normal because the preloaded certificate is self-signed and has the hostname [www.infoblox.com](http://www.infoblox.com), which does not match the destination IP address you entered in step 1. To stop the warning messages from occurring each time you log in to Grid Manager, you can generate a new self-signed certificate or import a third-party certificate with a common name that matches the FQDN (fully-qualified domain name) of the appliance. This is a very simple process. For information about certificates, see [Setting Login Options](#).
3. Enter the default username and password (**admin** and **infoblox**) on the Grid Manager login page, and then click **Login** or press Enter. For information, see [Logging on to the NIOS UI](#).
4. Read the Infoblox End-User License Agreement, and then click **I Accept** to proceed.
5. Read about the *Infoblox Customer Experience Improvement Program* and choose whether to participate (opt in) or not participate (opt out) in the program. By default, participation is enabled. If you want to opt out of the program, select **To Opt-Out of the alert program, please click here**. For more information about the program, see [Configuring the Customer Experience Improvement Program](#).
6. Click **OK**. Grid Manager may take a few seconds to load your user profile.
7. In the first screen of the *NIOS Setup* wizard, complete the following:
  - **Type of Network Connectivity:** Select **IPv4** as the communication protocol from the drop-down list.
  - Select **Configuring an HA pair** and click **Yes** to configure the first appliance.
  - **Send HA and Grid Communication over:** Select **IPv4** from the drop-down list for VRRP advertisements.
8. In the *NIOS Startup* wizard, select **Configuring an HA pair**. Click **Yes** to configure the first appliance.
9. Click **Next** and complete the following to configure network settings:
  - **Host Name:** Enter **ns3.corpxyz.com**.
  - **HA Pair Name:** Use the default name **Infoblox**.
  - **Shared Secret:** Enter **37eeT1d**.
10. Click **Next** and complete the following to set properties for the first node:
  - **Virtual Router ID:** Enter **150**.
  - **Required Ports and Addresses:** In the table, click the empty fields and enter the following information for each corresponding interface in the table:
    - VIP (IPv4): **10.1.4.10**

- Node 1 HA (IPv4): **10.1.4.7**
- Node 2 HA (IPv4): **10.1.4.9**
- Node 1 LAN1 (IPv4): **10.1.4.6**
- Node 2 LAN1 (IPv4): **10.1.4.8**
- Subnet Mask: **255.255.255.0**
- Gateway: **10.1.4.1**

Note that Some fields are prepopulated by Grid Manager based on the existing configuration of the appliance. All fields are required.

11. Click **Next** and complete the following to set admin password:
  - **Would you like to set admin password?:** Click **No**.
12. Click **Next** and complete the following to configure time settings:
  - **Time Zone:** Select UMT – 8:00 Pacific Time (US and Canada), Tijuana from the drop-down list.
  - **Would you like to enable NTP?:** Select **Yes** to synchronize the time with external NTP servers, and then click the Add icon. Grid Manager adds a row to the NTP Server table. Click the row and enter **10.120.3.10** in the **NTP Server** field.
13. Click **Next** to view the summary of the configuration. Review the information and verify that it is correct. You can change the information you entered by clicking **Previous** to go back to a previous step.
14. Click **Finish**.

## Node 2

1. From the **System** tab, select the **System Manager** tab, and then click **System Properties -> Setup Wizard** from the Toolbar.
2. In the first screen of the *NIOS Setup* wizard, complete the following:
  - **Type of Network Connectivity:** Select **IPv4** as the communication protocol from the drop-down list.
  - Select **Configuring an HA pair** and click **Yes** for configuring node 2 of the HA pair.
3. In the *NIOS Startup* wizard, select **Configuring an HA pair** to configure an independent HA pair. Click **No** for configuring node 2 of the HA pair.
4. Click **Next**, and then complete the following to configure network settings:
  - **HA Virtual IP address:** Enter **10.1.4.10**.
  - **HA Pair Name:** Use the default name **Infoblox**.
  - **Shared Secret:** Enter **37eeT1d**.
  - **Show Password:** Click this to display the shared secret.
5. Click **Next**, and then complete the following to set properties for the second appliance:
  - **IP Address:** Enter **10.1.4.8**.
  - **Subnet Mask:** Enter **255.255.255.0**.
  - **Gateway:** Enter **10.1.4.1**.
6. Click **Next** to view the summary of the configuration. Review the information and verify that it is correct. You can change the information you entered by clicking **Previous** to go back to a previous step.
7. Click **Finish**.

The setup of the HA pair is complete. From now on, when you make an HTTPS connection to the HA pair, use the VIP address 10.1.4.10.

## Enabling Zone Transfers

To allow the NIOS appliance to import zone data from the legacy server at 10.1.4.11, you must configure the legacy server to allow zone transfers to the appliance at 10.1.4.10.

## Legacy BIND Server

1. Open the named.conf file using a text editor and change the allow-transfer statement to allow zone transfers to the appliance at 10.1.4.10. For a sample of the required changes to the named.conf file.
2. After editing the named.conf file, restart DNS service for the change to take effect.

Legacy Windows 2000/2003 Server

Navigate to the *corpxyz.com Properties* dialog box, and then add 10.1.4.10 to the list of IP addresses to which you want to allow zone transfers.

## Importing Zone Data

You can import zone data from a legacy server to an independent HA pair. Use the following information:

- Forward-mapping zone: **corpxyz.com**
- Import zone from: **10.1.4.11**
- Reverse-mapping zone: **1.1.1.0**

## Defining Networks, Reverse-Mapping Zones, DHCP Ranges, and Infoblox Hosts

In this task, you enter data manually. For large data sets, you have the option of using the *Data Import Wizard* for loading DNS and DHCP configurations and data to make the process more efficient. To download the *Data Import Wizard*, visit [www.infoblox.com/import/](http://www.infoblox.com/import/).

## Networks

You can create all the subnetworks individually (which in this example are 10.1.1.0/24, 10.1.2.0/24, 10.1.4.0/24, and 10.1.5.0/24), or you can create a parent network (10.1.0.0/16) that encompasses all the subnetworks and then use the Infoblox split network feature to create the individual subnetworks automatically. The split network feature accomplishes this by using the IP addresses that exist in the forward-mapping zones to determine which subnets it needs to create. This example uses the split network feature. For information about creating networks, see [Configuring IPv4 Networks](#).

1. From the **Data Management** tab, select the **IPAM** tab, and then click **Add -> Add IPv4 Network** from the Toolbar.
2. In the *Add Network* wizard, complete the following:
  - **Address:** 10.1.0.0
  - **Netmask:** Use the netmask slider to select the /16 (255.255.0.0) netmask.
3. Click Next to select a server. Click the Add icon. Grid Manager displays ns3.corpxyz.com in the table.
4. Click **Save & Close**.
5. On the **IPAM** tab, select the **10.1.0.0/16** checkbox, and then select **Split** from the Toolbar.
6. In the *Split Network* dialog box, complete the following:
  - **Subnetworks:** Move the slider to 24.
  - **Immediately Add:** Select **Only networks with ranges and fixed addresses**.
  - **Automatically create reverse-mapping zones:** Select this checkbox.
7. Click **OK**.  
The appliance creates the following 24-bit subnets for the imported Infoblox hosts:
  - 10.1.1.0/24
  - 10.1.2.0/24
  - 10.1.4.0/24
  - 10.1.5.0/24
8. From the **IPAM** tab, select the 10.1.1.0/24 checkbox, and then click the Edit icon.
9. In the *DHCP Network* editor, enter information on the following tabs:
  - *General:* **Comment:** MGT
  - *Server Assignment:* Add ns3.corpxyz.com as a server
10. Click **Save & Close**.
11. To modify the other networks, repeat steps #8 – 10 for each network and use the following information:
  - 10.1.2.0/24 Network:
    - **Comment:** Dev
    - **Server Assignment:** ns3.corpxyz.com
  - 10.1.4.0/24 Network:
    - **Comment:** Server
    - **Server Assignment:** ns3.corpxyz.com
  - 10.1.5.0/24 Network:

- **Comment:** DMZ
- **Server Assignment:** ns3.corpxyz.com

## DHCP Ranges

1. On the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **10.1.1.0/24**, and then click **Add** -> **DHCP Range** from the Toolbar.
2. In the *Add Range* wizard, complete the following:  
**Start:** 10.1.1.10  
**End:** 10.1.1.50
3. Click **Next**, and then select **Server**. Grid Manager displays ns3.corpxyz.com as the assigned member.
4. Click **Save & Close**.
5. In the **Networks** tab, click **10.1.2.0/24**, and then click **Add** -> **DHCP Range** from the Toolbar.
6. In the *Add Range* wizard, complete the following:  
**Start:** 10.1.2.10  
**End:** 10.1.2.50
7. Click **Next**, and then select **Server**. Grid Manager displays ns3.corpxyz.com as the assigned member.
8. Click **Save & Close**.

## Infoblox Hosts

Defining both a MAC and IP address for an Infoblox host definition creates a DHCP host entry—like a fixed address—that you can manage through the host object. To add a MAC address to each host record that the appliance created when you imported forward—and reverse—mapping zone records:

1. On the **Data Management** tab, select the **IPAM** tab -> **10.1.1.0/24** -> **10.1.1.2**.
2. In the **Related Objects** tab, select the checkbox of the host record, and then click the Edit icon.
3. In the *Host Record* editor, click the MAC Address field, and then enter the following:
  - **MAC Address:** 00:00:00:aa:aa:aa
4. Click **Save & Close**.
5. Follow steps 1 – 4 to modify hosts with the following information:
  - printer2
    - IP Address: 10.1.2.2
    - MAC Address: 00:00:00:bb:bb:bb
  - storage1
    - IP Address: 10.1.4.2
    - MAC Address: 00:00:00:dd:dd:dd
  - storage2
    - IP Address: 10.1.4.3
    - MAC Address: 00:00:00:ee:ee:ee
  - proxymail
    - IP Address: 10.1.4.4
    - AC Address: 00:00:00:ff:ff:ff
  - proxyweb
    - IP Address: 10.1.4.5
    - MAC Address: 00:00:00:11:11:11
  - www
    - IP Address: 10.1.5.5
    - MAC Address: 00:00:00:55:55:55
  - mail
    - IP Address: 10.1.5.6
    - MAC Address: 00:00:00:66:66:66
  - ftp
    - IP Address: 10.1.5.7
    - MAC Address: 00:00:00:77:77:77

## Defining Multiple Forwarders

Since ns3.corpxyz.com is an internal DNS server, you configure it to forward DNS queries for external DNS name resolution to the primary and secondary DNS servers—ns1.corpxyz.com at 10.1.5.2 and ns2.corpxyz.com at 2.2.2.2.

1. From the **Data Management** tab, select the **DNS** tab, and then select **System DNS Properties** from the Toolbar.
2. In the *System DNS Properties* editor, click the Add icon on the **Forwarders** tab. Grid Manager adds a row to the table. Complete the following:
  - **Address:** Type **2.2.2.2**. Click **Add** again to add another forwarder.
  - **Address:** Type **10.1.5.2**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Each of the forwarders is assigned a random response time. The appliance sends the initial outbound query to the forwarder that has the lowest response time. If the first forwarder does not reply, the appliance tries the one with the next lowest random response time. The appliance adjusts and keeps track of the response times of the forwarders and uses the quicker one for future queries. If the quicker forwarder does not respond, the appliance then uses another one.

## Enabling Recursion on External DNS Servers

Since the HA pair forwards outbound queries to the two external DNS servers ns1.corpxyz.com (10.1.5.2) and ns2.corpxyz.com (2.2.2.2) for resolution, you must enable recursion on those servers. When a DNS server employs recursion, it queries other DNS servers for a domain name until it either receives the requested data or an error that the requested data cannot be found. It then reports the result back to the server that queried—in this case, the internal DNS server ns3.corpxyz.com (10.1.4.10), which in turn reports back to the DNS client.

### Infoblox Server in the DMZ Network (ns1.corpxyz.com, 10.1.5.2)

1. On the **Data Management** tab, select the **DNS** tab, and then click **System DNS Properties** from the Toolbar.
2. In the *System DNS Properties* editor, select the **Allow Recursion** checkbox on the **Queries** tab, and then click the Add icon -> **IPv4 Address**. Grid Manager adds a row to the **Allow recursive queries from** table. Complete the following:
  - **Permission:** Select **Allow** from the drop-down list.
  - **Name:** Enter **10.1.1.52**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

### BIND Server at ISP Site (ns2.corpxyz.com, 2.2.2.2)

1. Open the named.conf file using a text editor and change the recursion and allow-recursion statements to allow recursive queries from 1.1.1.8 (the NAT address of ns3).

```
options {
zone-statistics yes;
directory "/var/named/named_conf"; version"";
recursion yes;
listen-on { 127.0.0.1; 2.2.2.2; };
...
allow-recursion {1.1.1.8;};
transfer-format many-answers;
};
```

2. After editing the named.conf file, restart DNS service for the change to take effect.



Windows 2000/2003 Server at ISP Site (ns2.corpxyz.com, 2.2.2.2)

1. Click **Start -> All Programs -> Administrative Tools -> DNS**.
2. Right-click **ns3**, and then select **Properties -> Advanced**.
3. On the *Advanced* page in the *ns3 Properties* dialog box, clear the **Disable recursion** checkbox.
4. To save the configuration change and close the *ns3 Properties* dialog box, click **OK**.

## Modifying the Firewall and Router Configurations

Configure the firewall and router in your internal network to allow the following DHCP, DNS, and NTP traffic:

- To allow messages to pass from the DHCP clients in the DMZ—the web, mail, and FTP servers—to ns3 in the Server network, configure policies and DHCP relay agent settings on the firewall.
- To forward DHCP messages from DHCP clients in the MGT and Dev networks to ns3 in the Server network, configure relay agent settings on the router.
- To translate the private IP address of ns3 (10.1.4.10) to the public IP address (1.1.1.8) when forwarding DNS queries from ns3 to ns2, set a MIP (mapped IP) address on the firewall.
- To allow DNS queries from ns3 to ns1 and ns2 and NTP traffic from ns3 to the NTP server, configure firewall policies.

### Firewall

For example, enter the following commands on a Juniper firewall running ScreenOS 4.x or later:

#### DHCP Relay Configuration

```
set address trust ns3 10.1.4.10/32
set interface ethernet2 dhcp relay server-name 10.1.4.10
set policy from dmz to trust ns1 ns3 DHCP-Relay permit
```

#### DNS Forwarding

```
set interface ethernet1 mip 1.1.1.8 host 10.1.4.10
set policy from trust to untrust ns3 ns2 dns permit
set policy from trust to dmz ns3 ns1 dns permit
```

#### NTP

```
set policy from dmz to untrust ns1 ntp_server ntp permit
```

### Router

For example, enter the following commands on a Cisco router running IOS for release 12.x or later:

#### DHCP Relay Configuration

```
interface ethernet1
ip helper-address 10.1.4.10 interface ethernet2
```

```
ip helper-address 10.1.4.10
```

## Enabling DHCP and Switching Service to the NIOS Appliance

With the Infoblox in place and the firewall and router configured for relaying DHCP messages, you can switch DHCP service from the legacy DHCP server at 10.1.4.11 to the HA pair at 10.1.4.10 (VIP address).



### Note

To minimize the chance of duplicate IP address assignments during the transition from the legacy DHCP server to the appliance, shorten all lease times to a one-hour length in advance of the DHCP server switch. Then, when you take the legacy DHCP server offline, the DHCP clients quickly move to the new server when their lease renewal efforts fail, and they broadcast DHCPDISCOVER messages. To determine how far in advance you need to shorten the lease length, find the longest lease time (for example, it might be two days). Then change the lease length to one hour at a slightly greater interval of time before you plan to switch DNS service to the appliance (for example, three days before the switch over).

By changing the lease length this far in advance, you can be sure that all DHCP leases will be one-hour leases at the time of the switch-over. If the longest lease length is longer—such as five days—and you want to avoid the increased amount of traffic caused by more frequent lease renewals over a six-day period, you can also employ a stepped approach: Six days before the switch-over, change the lease lengths to one-day leases. Then two days before the switch-over, change them to one-hour leases.

1. Open an Internet browser window, enter <https://10.1.4.10>, and then log in to the appliance using the username **admin** and password **SnD34n534**.
2. From the **Data Management** tab, select the **DHCP** tab, and then click **Start** from the Toolbar.
3. In the *Start Member DHCP Service* dialog box, click **Yes**. The HA pair is ready to provide DHCP service to the network.
4. Take the legacy DHCP server at 10.1.4.11 offline.  
When the DHCP clients are unable to renew their leases from the legacy DHCP server, they broadcast DHCPDISCOVER messages to which the new DHCP server responds.

## Managing and Monitoring

Infoblox provides tools for managing IP address usage and several types of logs to view events of interest and DHCP and DNS data. After configuring the appliance, you can use the following resources to manage and monitor IP address usage, DNS and DHCP data, and administrator and appliance activity.

### IPAM (IP Address Management)

IPAM offers the following services:

- **Simple IP address modification:** Within a single IP address-centric data set, you can modify the Infoblox host, DHCP, and DNS settings associated with that IP address.
- **Address type conversion:** Through IPAM functionality, you can make the following conversions:
  - Currently active dynamic addresses to fixed addresses, reserved addresses, or Infoblox hosts.
  - Fixed addresses to reservations or hosts.
  - Reservations to hosts.
- **Device classification:** You can make detailed descriptions of appliances in DHCP ranges and appliances defined as Infoblox hosts and as fixed addresses.
- **Three distinct views of IP address usage:** To monitor the usage of IP addresses on your network, you can see the following different views:
  - High-level overall network view: On the **Data Management** tab, select the **IPAM** tab -> *member*. You can view the network usage in the Net Map or List view. You can also drill down to specific IP address to get detailed information.

- DHCP lease history records: From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Lease History**.

## Logs

The following list has some useful information:

- Logs, as described in [Monitoring the Appliance](#).
  - Audit Log – Contains administrator-initiated events.
  - System Log – Contains events related to hardware and software operations.
- DNS statistics, as described in [Configuring DNS Services](#).
  - DNS Configuration – Contains DNS server settings for the Infoblox DNS server.
  - Zone Statistics – Contains the results of all DNS queries per zone.
- DHCP information, as described in [Configuring DHCP Properties](#).
  - DHCP Configuration – Contains DHCP server settings and network, DHCP range, and host settings for the Infoblox DHCP server.
  - DHCP Leases – Contains a real-time record of DHCP leases.
  - DHCP Lease History – Contains a historical record of DHCP leases.
  - DHCP Statistics – Contains the number of currently assigned static and dynamic addresses, and the high and low watermarks per network.
  - Network Statistics – Contains the number of static hosts, dynamic hosts, and available hosts per network.

## Verifying the Deployment

After you deploy a single independent appliance or HA pair, you can make an HTTPS connection to it, log in, and check its status.

### Single Independent Appliance

From the Dashboard, check the appliance status in the *System Status* widget. For information, see [Member Status \(System Status\)](#).

- If the Status icon is green, the appliance has a network connection and is operating properly.
- If the Status icon is red, there is a problem. To determine what it is, look at the system log file for this appliance by selecting the **Administration** tab -> **Logs** tab -> **Syslog**.

### Independent HA Pair

1. Make an HTTPS connection to the VIP address of the HA pair, log in, and check the status of both nodes.
2. From the Dashboard, check the appliance status in the *SystemStatus* widget. For information, see [Member Status \(System Status\)](#).
  - If the Status icon is green, both nodes have connectivity with each other and are operating properly.
  - If the Status icon is yellow, the two nodes are in the process of forming an HA pair.
  - If the Status icon is red, the passive node is offline or there is a problem. To determine what it is, look at the system log file by selecting the **Administration** tab -> **Logs** tab -> **Syslog**. You can also gather information from the **System** tab -> **SystemManager** tab.

## Infoblox Tools for Migrating Bulk Data

Typically, the next step after cabling a single independent appliance to a network and configuring its network settings—or cabling two independent appliances to a network and configuring them as an HA pair—is to import data from legacy DNS, DHCP, and TFTP servers. Infoblox provides several tools to accomplish this:

- The CSV import feature allows you to import DNS, DHCP, and IPAM data through Grid Manager. You can add, overwrite, or merge data using this feature. The appliance updates the database based on import settings and the data you specify in the data files. From the **Data Management** tab of Grid Manager, you can access the *Import*

*Manager* editor from which you start a data import. You can also export existing data to a CSV file. You can use this file to modify data, and then re-import the data into the database using the CSV import feature. For information, see [About CSV Import](#).

- The Infoblox Data Import Wizard is a useful tool that simplifies the importation of DNS, DHCP and IPAM, and TFTP settings and data into a NIOS appliance. For large data sets, this option is an efficient approach. To download the Data Import Wizard, visit [www.infoblox.com/import/](http://www.infoblox.com/import/).
- For smaller DNS data sets, you can use the zone import feature, which allows you to import data on a per-zone basis (see [Importing Zone Data](#)).

## Deploying Cloud Network Automation

The Infoblox Cloud Network Automation solution automates IPAM (IP address management) for physical and virtual network devices on your CMP (whether it is private, public, or hybrid). Instead of manually provisioning IP addresses and DNS name spaces for network devices and interfaces, you can use Cloud Network Automation to leverage DNS and DHCP features of the Grid to manage your cloud networks. When your cloud consists of a large number of servers and VMs (virtual machines) that have multiple associated network interfaces, manually provisioning and de-provisioning IP addresses and managing DNS and DHCP data can become error-prone. Utilizing Cloud Network Automation can minimize human errors by streamlining IPAM, improve visibility of your cloud networks, and maximize the flexibility and efficiency that virtualization offers in your cloud environment.

To maximize the capability of Cloud Network Automation, ensure that you understand the operations of Cloud Network Automation and how to deploy necessary components to suit your network topology and CMP. For detailed information about the Cloud Network Automation operations, see [Cloud Network Automation](#).

With valid licenses installed, you can configure Cloud Platform Appliances in your Grid to provide DNS and DHCP services for virtual machines and network devices deployed through a CMP (Cloud Management Platform). To deploy a Grid that contains cloud members, it is important to understand the required components, how to configure the members to enable Cloud Network Automation, and which cloud API objects are supported.

- [Cloud Network Automation](#)
- [Licensing Requirements](#)
- [Licensing Configurations](#)
- [Administrative Permissions](#)
- [Cloud Network Automation Administrative Permissions](#)
- [Supported Cloud Platform Appliance Models](#)
- [Setting Up Cloud Network Automation](#)
- [About Cloud API Service](#)
- [About Cloud API Requests](#)
- [About Authority Delegation](#)
- [Configuring Grid and Member Cloud API Properties](#)
- [Extensible Attributes for Cloud Objects](#)
- [Viewing Cloud Objects](#)

### Related topic

[Cloud Network Automation](#)

## Cloud Network Automation

The Cloud Network Automation figure shows the basic concept and operations of Cloud Network Automation, which includes two major components: the Grid Master that has a Cloud Network Automation license installed and one or more Cloud Platform Appliances that provide the ability to process API requests. Instead of sending all API requests to the Grid Master, you can send requests directly to the Cloud Platform Appliances. The Cloud Network Automation license installed on the Grid Master enables visibility and reporting on cloud tenants, VM IP addresses, and DNS record allocation. This license can be used in conjunction with Cloud Platform Appliances to provide local survivability and additional scalability of cloud API requests within individual data centers, or it can be used with an existing Grid Master servicing all cloud API

requests.

A Cloud Platform Appliance is a Grid member designed and dedicated to accept and process WAPI (RESTful API) requests related to cloud objects, in addition to serving DNS and DHCP protocols. You can deploy multiple Cloud Platform Appliances within your Grid to scale the processing of API requests or to provide redundancy. Cloud Platform Appliances include built-in HTTPS proxy capability that redirects cloud API requests to the appropriate Cloud Platform Appliance or to the Grid Master for processing. In other words, cloud API requests can be sent to any of the Cloud Platform Appliances within the Grid and the call is either processed locally or transparently forwarded to the appliance that is authoritative for the object referenced in the cloud API request. For information about supported Cloud Platform Appliances and their specifications, see [Supported Cloud Platform Appliance Models](#). Once you have installed the Cloud Platform license on the appliance, fixed address configuration takes effect immediately by default and no DHCP service restart is required on the Cloud Platform Appliance. For information about this feature, see [Configuring Fixed Addresses without Restarting DHCP Service](#). You can also add and delete IPv4 and IPv6 fixed addresses, reservations, and host records within any delegated IP address ranges through Grid Manager (the Infoblox GUI), in addition to using cloud API calls. For more information, see [Managing IPv4 DHCP Data](#) and [Managing IPv6 DHCP Data](#).

On the CMP, you can either deploy a cloud adapter and use it as the cloud API client for sending cloud API requests to the Cloud Platform Appliances, or you can customize your CMP to make cloud API requests directly to Cloud Platform Appliances or to the Grid Master. The cloud adapter can be configured to send API requests always to a single Cloud Platform Appliance or to different Cloud Platform Appliances to handle situations where the primary Cloud Platform Appliance may not be available or to distribute API load among multiple Cloud Platform Appliances. Infoblox Cloud Network Automation supports the following cloud adapters: Infoblox IPAM Plug-In for VMware, OpenStack Adapter, and AWS (Amazon Web Services) API Proxy. For information about the IPAM Plug-In for VMware and OpenStack Adapter, refer to their respective Quick Start Guides. For information about the AWS API Proxy and how to set up AWS configurations, refer to the [Infoblox Installation Guide for NIOS for AWS](#).

In order to distribute API processing and provide additional scalability both for updating APIs and serving DNS and DHCP protocols, Cloud Network Automation enables you to delegate specific sets of IPAM, DNS, and DHCP data to one or more Cloud Platform Appliances. Once authority for an object or set of objects has been delegated to a Cloud Platform Appliance, cloud API requests to create, modify, or delete objects under the scope of delegation for that appliance are processed locally and available immediately for serving DNS and DHCP to VMs within the cloud. This eliminates the need to send requests to the Grid Master to create, modify, or delete objects within the Grid. Changes made to objects on individual Cloud Platform Appliances are synchronized with the Grid Master in near real time using Grid replication to provide centralized visibility while retaining distributed processing capability. If a Cloud Platform Appliance is not authoritative for the object referenced in the cloud API requests, it automatically proxies that request to the Cloud Platform Appliance that is authoritative for the object or to the Grid Master (if it is authoritative for the object). Similarly, cloud API requests made to the Grid Master are proxied to the authoritative Cloud Platform Appliance or processed locally on the Grid Master if it is authoritative for the object. For information about authority delegation for supported objects, see [About Authority Delegation](#). For information about proxying cloud API requests, see [About Cloud API Requests](#).

Cloud API requests are processed through the cloud API service that operates on the Cloud Platform Appliance. This service can also be enabled on the Grid Master as well as other Cloud Platform Appliances. The cloud API service is HTTPS-based; therefore, to ensure that the cloud API service functions properly, port 443 for HTTPS connectivity must be open between the CMP and each Cloud Platform Appliance and/or the Grid Master receiving the cloud API requests. To ensure that the proxying function works properly, port 443 for HTTPS must be open bi-directionally between each of the Cloud Platform Appliances as well as between each Cloud Platform Appliance and the Grid Master. You must also configure your firewalls and ACLs accordingly. Note that this service uses the VIP address on each Infoblox appliance as the destination address.

All objects created, modified, or deleted by the cloud adapter are reflected in the NIOS database. You can view cloud objects and their associated data in the **Cloud** tab of Grid Manager if the Cloud Network Automation license is installed on the Grid Master. For more information, see [Viewing Cloud Objects](#). Note that it is possible to use Cloud Platform Appliances without deploying the Cloud Network Automation license. However, without the Cloud Network Automation license, VM and tenant information is only displayed as extensible attributes associated with IPAM, DHCP, and DNS objects in Grid Manager rather than in separate tables under the **Cloud** Tab.

Before you can send cloud API requests to a Cloud Platform Appliance or the Grid Master, you must create admin groups that have cloud API access. Only admin users that have cloud API access and applicable permissions may be used for sending cloud API requests. If the Cloud Network Automation license is installed on the Grid Master, it is also possible to assign **Tenant** permissions to admin users to restrict these users to only be able to view objects related to a given tenant or a set of tenants. For information about admin groups and how to manage admin users, see [Managing Administrators](#).

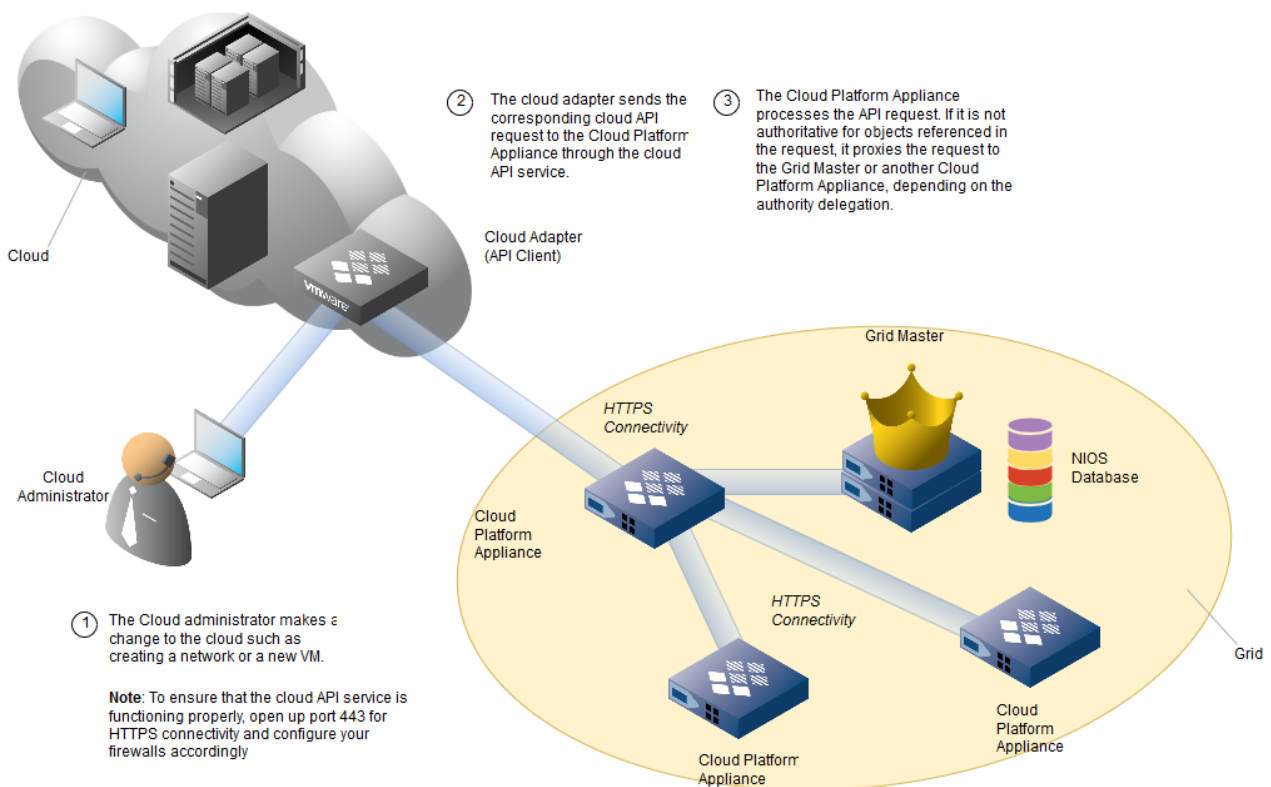
Note that there is no current capability to bi-directionally synchronize NIOS data with CMP data. Therefore, cloud information in NIOS is accurate only up to the point when specific cloud API requests are received by the Cloud Platform Appliance from an adapter running on a CMP. Only cloud information obtained through API requests to the Infoblox API service will be available in NIOS.



**Note**

Unlike standard WAPI requests, all cloud API related events are logged to the NIOS syslog instead of the NIOS audit log.

*Cloud Network Automation*



**Related topic**

[Deploying Cloud Network Automation](#)

**Licensing Requirements**

To enable Cloud Network Automation, you must install valid licenses on the Grid Master and Cloud Platform Appliance members. Depending on your deployment scenarios, you can take advantage of Elastic Scaling to automatically deploy virtual cloud appliances, either inside or outside your CMP.

The following valid licenses are part of the Cloud Network Automation solution:

- Cloud Network Automation license on the Grid Master and Grid Master Candidate
  - Cloud Platform license on the Cloud Platform Appliances
- The license you install on the Grid Master enables the **Cloud** user interface functions in Grid Manager and **Tenant** permissions.



The license you install on the Cloud Platform Appliance enables the cloud API service on the Cloud Platform Appliance. Note that it is possible to use the Cloud Network Automation solution only with the Cloud Network Automation license or with one or more Cloud Platform Appliances. In the case when only the Cloud Network Automation license is installed on the Grid Master, all cloud API requests are sent to the Grid Master instead of to individual Grid members. Creation of cloud objects through cloud API requests is visible in the **Cloud** tab of Grid Manager on the Grid Master.

When Cloud Platform Appliances are used without the Cloud Network Automation license, cloud API requests are sent either to the Cloud Platform Appliances or to the Grid Master. However, the **Cloud** tab in Grid Manager is not available on the Grid Master for viewing cloud objects created through cloud API requests.

You can also use the CLI command `set temp_license` to generate and install temporary licenses. This provides licensed features and functionality for the interim, while you wait for your permanent licenses to arrive. For information about how to install a temporary license, see [Adding Temporary Licenses](#). Note that the temporary license is only effective on the Grid Master, not the Grid Master Candidate.

For an HA pair, ensure that you use the same appliance models for both nodes and install the Cloud Platform license on both nodes as well. For information about supported models, see [Supported Cloud Platform Appliance Models](#). If a failover occurs and the passive node does not have a valid license, the cloud API service will stop and all resource delegations to the Cloud Platform Appliance will also stop.

Note that the Cloud Network Automation license on the Grid Master is incompatible with the following licenses:

- Multi-Grid Manager
- Reporting

Cloud Platform licenses are only supported on Cloud Platform Appliances. They may not be installed on any other Infoblox physical or virtual appliances. The following licenses and functionality are not supported on the Cloud Platform Appliances:

- Microsoft Management
- Multi-Grid Management
- Network Insight
- Reporting
- Tiered DNS Cache Acceleration
- DNS Cache Acceleration
- Load Balancing
- Infoblox
- RIR (Regional Internet Registry)

Before you install or remove the Cloud Platform license, consider the following:

- Installing or removing the Cloud Platform license stops the cloud API service.
- When you remove the Cloud Platform license from the appliance, it still serves DNS and DHCP if those licenses are installed on the appliance. However, the appliance will no longer be able to run the cloud API service. In addition, you cannot delegate authority to this member for objects that have not already been delegated to this appliance. Existing delegations to this member remain in the NIOS database, but API requests proxied from other Cloud Platform Appliances or from the Grid Master will fail.

## Licensing Configurations

Depending on your Grid configuration and how you want to deploy Cloud Network Automation, Infoblox supports the following licensing configurations. For information about cloud licenses and how to install them, see [Licensing Requirements](#).

- **In an Infoblox Grid, the Grid Master has the Cloud Network Automation license installed and the Cloud Platform Appliance has the Cloud Platform license installed:** The Grid Master can process both regular RESTful API and cloud API requests and the Cloud Platform Appliance can process cloud API requests. With the Cloud Platform license installed, cloud API requests can be proxied among the Grid Master and Cloud Network Appliances based on the delegation authority of the referenced objects. You can also manage the cloud API service and cloud objects through the **Cloud** tab in Grid Manager
- **In an Infoblox Grid, the Grid Master does not have the Cloud Network Automation license installed but the Cloud Platform Appliance has the Cloud Platform license installed:** The Grid Master can process both regular RESTful API and cloud API requests and the Cloud Platform Appliance can process cloud API requests. With the Cloud Platform license installed, cloud API requests can be proxied among the Grid Master and Cloud Network

Appliances based on the authority delegation of the referenced objects. Without the Cloud Network Automation license however, only objects whose authority has been delegated to Cloud Platform Appliances can be managed through Grid Manager. You will not have visibility of cloud objects, such as tenant information, through Grid Manager because the **Cloud** user interface function (the **Cloud** tab) is not available without the Cloud Network Automation license. You may manage cloud objects and their respective data through cloud API requests.

- **In an Infoblox Grid with other Grid members but without Cloud Platform Appliances, the Grid Master has the Cloud Network Automation license installed:** The Grid Master can process both regular RESTful API and cloud API requests. You can also manage cloud objects through the **Cloud** tab in Grid Manager.
- **Standalone Grid Master with the Cloud Network Automation license installed:** The appliance can process both regular RESTful API and cloud API requests. You can also manage the cloud API service and cloud objects through the **Cloud** tab in Grid Manager.
- **Standalone Grid Master without the Cloud Network Automation license installed:** The appliance can process only regular RESTful API requests, but not cloud API requests. You cannot manage the cloud API service nor cloud objects through Grid Manager because the **Cloud** user interface function (the **Cloud** tab) is not available.

The following table summarizes the supported licensing scenarios:

	Cloud Network Automation License on Grid Master	Cloud Platform License on Cloud Platform Appliance	Supports Regular RESTful API Calls on Grid Master	Supports Cloud API Calls on Grid Master and/or Cloud Platform Appliance(s)	Cloud User Interface Function in Grid Manager
Infoblox Grid <b>with</b> Cloud Platform Appliance (s)	✓	✓	✓	✓	✓
Infoblox Grid <b>with</b> Cloud Platform Appliance (s)	✗	✓	✓	✓	✗
Infoblox Grid with other Grid members but <b>without</b> Cloud Platform Appliance (s)	✓	N/A	✓	✓	✓
Standalone Grid Master	✓	N/A	✓	✓	✓
Standalone Grid Master	✗	N/A	✓	✗	✗

## Administrative Permissions

You must define admin users and their permissions in the admin group and assign specific roles to it before you can use these admin users to send cloud API requests. You can also define object permissions to specific admin groups or admin



users so they can manage specific objects through cloud API requests. For more information, see [About Admin Accounts](#) and [About Admin Groups](#).



#### Note

When you deploy Cloud Network Automation, the *cloud-api-only* is created automatically. You cannot delete this admin group.

Depending on where a cloud API request is sent and whether the scope of delegation for an object is explicit or implicit, permissions configured for the admin user and object may or may not apply. In addition, depending on the objects referenced in cloud API requests, specific restrictions may apply. For supported objects and their restrictions, see [Supported Cloud API Objects](#).

For cloud API requests, admin permissions are applied based on the delegation status of the objects referenced in the requests. If an object is not delegated (owned by the Grid Master) and the cloud API request is sent directly to the Grid Master or proxied to the Grid Master, all applicable admin and object permissions apply. On the other hand, if authority for an object referenced in a cloud API request is explicitly delegated to a Cloud Platform Appliance and the request is sent to this appliance, the admin user has full permission for this object within the scope of delegation. In this case, specific permissions configured for the admin user and the referenced object are ignored. For more information about admin and object permissions, see [About Administrative Permissions](#).

It is important to note that once you delegate authority of an object to the Cloud Platform Appliance, specific admin and object permissions are not enforced. Therefore, if you do not want certain objects to be created or modified through a cloud API request, do not delegate the authority of these objects and their parent objects to a Cloud Platform Appliance. For example, if you do not want host records to be created through cloud API requests, do not delegate the authority of the relevant networks, zones, or both to the Cloud Platform Appliance. On the other hand, if you want the ability to restrict permissions for specific objects referenced in cloud API updates, you can create different admin groups or admin users that are authorized to make cloud API updates on respective Cloud Platform Appliances. The following example illustrates this capability.

### Configuration Example

If you want to restrict the creation and modification of records for networks 10.10.10.0/24 and 10.10.20.0/24 through cloud API updates, do the following:

1. Create two admin users **APIUser1** and **APIUser2** in an admin group.
2. Delegate the authority of network 10.10.10.0/24 to Cloud Platform Appliance 1 (**CP1**) and 10.10.20.0/24 to Cloud Platform Appliance 2 (**CP2**).
3. On **CP1**, add **APIUser1** and on **CP2**, add **APIUser2** to the list of administrators that can send cloud API requests, as described in [Configuring Grid and Member Cloud API Properties](#).

Now when you use **APIUser1** to send cloud API requests, you can add and modify records for network 10.10.10.0/24, but you cannot do so for network 10.10.20.0/24. Conversely, you can add and modify records for network 10.10.20.0/24 only when you use **APIUser2**.

## Cloud Network Automation Administrative Permissions

You must define admin users and their permissions in the admin group and assign specific roles to it before you can use these admin users to send cloud API requests. You can also define object permissions to specific admin groups or admin users so they can manage specific objects through cloud API requests. For more information, see [About Admin Accounts](#) and [About Admin Groups](#).



#### Note

When you deploy Cloud Network Automation, the *cloud-api-only* is created automatically. You cannot delete this admin group.

Depending on where a cloud API request is sent and whether the scope of delegation for an object is explicit or implicit, permissions configured for the admin user and object may or may not apply. In addition, depending on the objects referenced in cloud API requests, specific restrictions may apply. For supported objects and their restrictions, see [Supported Cloud API Objects](#).

For cloud API requests, admin permissions are applied based on the delegation status of the objects referenced in the requests. If an object is not delegated (owned by the Grid Master) and the cloud API request is sent directly to the Grid Master or proxied to the Grid Master, all applicable admin and object permissions apply. On the other hand, if authority for an object referenced in a cloud API request is explicitly delegated to a Cloud Platform Appliance and the request is sent to this appliance, the admin user has full permission for this object within the scope of delegation. In this case, specific permissions configured for the admin user and the referenced object are ignored. For more information about admin and object permissions, see [About Administrative Permissions](#).

It is important to note that once you delegate authority of an object to the Cloud Platform Appliance, specific admin and object permissions are not enforced. Therefore, if you do not want certain objects to be created or modified through a cloud API request, do not delegate the authority of these objects and their parent objects to a Cloud Platform Appliance. For example, if you do not want host records to be created through cloud API requests, do not delegate the authority of the relevant networks, zones, or both to the Cloud Platform Appliance. On the other hand, if you want the ability to restrict permissions for specific objects referenced in cloud API updates, you can create different admin groups or admin users that are authorized to make cloud API updates on respective Cloud Platform Appliances. The following example illustrates this capability.

### Configuration Example

If you want to restrict the creation and modification of records for networks 10.10.10.0/24 and 10.10.20.0/24 through cloud API updates, do the following:

1. Create two admin users **APIUser1** and **APIUser2** in an admin group.
2. Delegate the authority of network 10.10.10.0/24 to Cloud Platform Appliance 1 (**CP1**) and 10.10.20.0/24 to Cloud Platform Appliance 2 (**CP2**).
3. On **CP1**, add **APIUser1** and on **CP2**, add **APIUser2** to the list of administrators that can send cloud API requests, as described in [Configuring Grid and Member Cloud API Properties](#).

Now when you use **APIUser1** to send cloud API requests, you can add and modify records for network 10.10.10.0/24, but you cannot do so for network 10.10.20.0/24. Conversely, you can add and modify records for network 10.10.20.0/24 only when you use **APIUser2**.

### Supported Cloud Platform Appliance Models

The following table lists the Infoblox vNIOS virtual appliances that you can use as Cloud Platform Appliances. A Cloud Platform Appliance can only be configured as a Grid member, not a standalone appliance. To configure an HA pair, you must use the same vNIOS models for both the active and passive nodes. For information about HA pairs, see [About HA Pairs](#).



#### Note

Cloud Platform Appliances do not support auto-provisioning. Pre-provisioning is supported for DNS and DHCP data. For information about pre-provisioning Cloud Platform Appliances, see [Pre-Provisioning NIOS and vNIOS Appliances](#).

#### Supported Cloud Platform Appliances

vNIOS Appliance	Storage (GB)	# of CPU Cores	Virtual CPU Core Frequency	Memory Allocation
CP-V805	250	2	2000 MHz	16 GB

CP-V1405	250	4	6000 MHz	32 GB
CP-V2205	250	8	12000 MHz	64 GB

## Setting Up Cloud Network Automation

To set up Cloud Network Automation, ensure that you understand the required components, licenses, and firewall configuration, as described in [Cloud Network Automation Operations](#). You also need to understand what a Grid is and how to set up a Grid and deploy vNIOS appliances. For more information, see [About Grids](#).

Following are high-level steps for setting up Cloud Network Automation:

1. Evaluate your network topology and IPAM requirements, and then decide how you want to set up your Grid in the cloud environment. Based on your requirements, install and configure a cloud adapter that supports Cloud Platform Appliances on your CMP. This adapter functions as the cloud API client. For information about how to install and configure the cloud adapter, refer to the *Quick Start Guide* for the cloud adapter you are deploying, available on the Infoblox Support site. If you are implementing Cloud Network Automation in your Amazon VPCs (Virtual Private Clouds), refer to the *Infoblox Installation Guide for vNIOS for AWS*.
2. Obtain valid licenses through your Infoblox representatives. Note that Cloud Platform licenses may only be installed on cloud virtual appliances. They may not be installed on physical or virtual Trinzic Enterprise appliances. Deploy the Cloud Platform Appliances and install the Cloud Platform license along with the Enterprise, vNIOS, DHCP, and DNS licenses included in the license bundle for the Cloud Platform Appliance. If you plan to use the Cloud user interface functionality on the Grid Master, obtain a Cloud Network Automation license for your Grid Master and Grid Master Candidate. For information about how to obtain and install licenses, see [Managing Licenses](#).

If you are implementing Cloud Network Automation in AWS, consider Elastic Scaling to see if it better fits your business needs when you purchase your licenses. You can purchase multiple dynamic licenses and automatically pre-provision vNIOS cloud appliances according to your business needs.

1. If the Cloud Platform Appliance is not part of the Grid, join it to the Grid. For more information about how to join a member to the Grid, see [Joining Appliances to the Grid](#). If you have implemented Elastic Scaling, ensure that you follow the necessary steps to automatically provision your vNIOS cloud appliances and join them to the Grid. For more information, see [Using Elastic Scaling to Pre-provision and Launch vNIOS Members](#).
2. Log in to Grid Manager and complete the following:
  - Start the cloud API service on the Cloud Platform Appliance, as described in [Starting the Cloud API Service](#). The cloud API service is disabled by default.
  - Define admin groups and admin permissions for Cloud Network Automation, as described in [About Admin Accounts](#) and [About Admin Groups](#).
  - Control which users can send API requests, as described in [Configuring Grid and Member Cloud API](#).
  - Add global and object permissions for admin users who can manage cloud objects, as described in [Managing Permissions](#).

You can now start sending API requests through the cloud adapter or AWS API Proxy. To understand the API request process and the supported API objects, see [About Cloud API Requests](#).

## About Cloud API Service

The cloud API service provides the ability to automate management of IP addresses and DNS records so your cloud environment can take full advantage of IPAM, DNS, and DHCP capabilities in NIOS without the need for manual intervention. This service is supported for the following scenarios:

- Communication between the cloud adapter (acting as an API client) and the Cloud Platform Appliance or between Cloud Platform Appliances. This cloud API service accepts and processes a subset of the WAPI requests that are currently supported on the Grid Master either directly from an adapter or proxied through another Cloud Platform Appliance or from the Grid Master.

- Communication between the cloud adapter and the Grid Master, or between Cloud Platform Appliances and the Grid Master. This cloud API service processes requests received directly from the cloud adapter or processes requests received by other Cloud Platform Grid members.

The admin users that you use to send cloud API requests must have applicable access to the cloud API in order for the API requests to be processed. For information about admin groups, see [Managing Admin Groups and Admin Roles](#).

## Starting the Cloud API Service

To start the cloud API Service:

1. From the **Grid** tab, select the **Services** tab -> *cloud\_member* checkbox.
2. Click **Cloud API** on the top navigation bar, and then click **Start** from the Toolbar.

## Monitoring Cloud API Service

To monitor the status of the cloud API service, from the **Cloud** tab, select the **Members** tab -> *cloud\_member* link. Grid Manager displays the service status, as described in [Service Status](#).

You can also monitor the cloud API service through the following:

- View the cloud API service through the Member service status, as described in [Member Status](#).
- Configure the appliance to receive SNMP traps for the cloud API service, as described in [Monitoring with SNMP](#).
- View *Cloud Statistics* widget on the Dashboard, as described in [Cloud Statistics](#).
- Review event messages in the syslog, as described in [Viewing the Syslog](#).

## About Cloud API Requests

In your cloud environment, the cloud adapter acts as the cloud API client. Only API requests made by admin users who have the correct permissions on the cloud API ACL (Access Control List) are processed by the cloud API service. When the Cloud Platform Appliance receives a cloud API request, it processes the request based on authority delegation of the objects and respective cloud extensible attributes. For information about cloud extensible attributes, see [Extensible Attributes for Cloud Objects](#). If the Cloud Platform Appliance is not authoritative for the referenced objects, it proxies the request to the authoritative appliance that can be another Cloud Platform Appliance or to the Grid Master if no authority delegation is defined. For information about proxying cloud API requests, see [Proxying Cloud API Requests](#) below.



### Note

For the cloud API service to function properly, configure your networks and firewalls accordingly to allow port 443 HTTPS connectivity between the cloud adapter and Cloud Platform Appliance, between the cloud adapter and the Grid Master (if applicable), between the Grid Master and Cloud Platform members, and between each Cloud Platform member.

If you are using the AWS API Proxy to send API requests, ensure that you understand how to set up and configure the proxy. For detailed information, refer to the [Infoblox Installation Guide for vNIOs for AWS](#).

When implementing Cloud Network Automation in AWS, you can use Elastic Scaling to allocate and deallocate dynamic licenses and automatically spin up vNIOs Grid members and join them to the Grid. You can purchase and install NIOs feature licenses in advance and store them in a license pool container on the Grid Master. You can then decide when and how to automatically provision and configure vNIOs for AWS cloud virtual appliances. When you remove a vNIOs cloud appliance, the licenses on this appliance are released and returned to the license pool and are available for the next deployment.

## Cloud API Request Process

As described in the below table, all cloud API requests are subject to the following process before responses are returned. The appliance

### *Cloud API Request Process*

Steps	Descriptions	Configuration that affects the outcome of this step
Authentication and categorization	All cloud API requests are authenticated based on the authentication sources. Once authenticated, the requests are categorized as either a cloud API request or not. All requests that specify user identity as users defined in admin groups with cloud API access are <a href="#">Managing Admin Groups and Admin Roles</a> categorized as cloud API requests.	Define admin user accounts that can be used to send cloud API requests.  For information, see <a href="#">Managing Admin Groups and Admin Roles</a> .
Authorization	All cloud API requests are subject to authorization based on the ACLs (Access Control Lists) defined for the Grid or Cloud Platform Appliance. You can control which admin accounts can be used to send API requests. The ACLs can contain admin users in admin groups with cloud API access or remote authenticated users.	Define ACLs on the Grid Master or Cloud Platform Appliance.  For information, see <a href="#">Configuring Grid and Member Cloud API Properties</a> .
Proxying Requests	If a Cloud Platform Appliance is not authoritative for a cloud API request, it proxies the request either to the authoritative Cloud Platform Appliance or to the Grid Master for processing. Similarly, if an object has been delegated and the API request is made to the Grid Master, the Grid Master proxies that request to the authoritative Cloud Platform Appliance.	Ensure that HTTPS connectivity between each Cloud Platform member and between each Cloud Platform member and the Grid Master is functioning properly for proxying.  For information, see <a href="#">Proxying Cloud API Requests</a> below.
Validation	NIOS performs a final validation on the cloud API request based on permissions configured for the admin users and restrictions for the applicable objects. If the request is processed within the scope of an explicit delegation, the admin user is considered to have full permissions within the scope, and any permission defined for admin groups with cloud API access is ignored. Otherwise, the request is subject to validation for all permissions defined for admin groups with cloud API access.	Define admin permissions for admin groups with cloud API access.  For information, see <a href="#">About Admin Groups</a> .
Auditing	Cloud API related events are logged to the NIOS syslog of the Grid member that processes the API requests instead of to the NIOS audit log.	Configure syslog server for the cloud member.  For information, see <a href="#">Viewing the Syslog</a> .

## Supported Cloud API Objects

The following table lists all the supported cloud API object types, methods, and functions. In your cloud API requests, you cannot include RESTful API object types, methods, and functions that are not listed in the table, even when the Grid Master supports them for other purposes. Note that the supported types and operations for cloud API requests are subsets of all types and operations supported on the Grid Master.

Before you send any cloud API requests, ensure that you understand the implications and restrictions for each supported

object. NIOS uses extensible attributes to associate specific information with a cloud object. For information about the default cloud extensible attributes and how to use them, see [Extensible Attributes for Cloud Objects](#). In AWS (Amazon Web Services), you can create a VPC (Virtual Private Cloud) and a subnet using the same network address and subnet mask. For example, you can add 172.29.02.0/24 as the VPC and 172.29.2.0/24 as the subnet and create VMs in the subnet. However, you cannot add a network container and a network using the same network address and subnet mask in NIOS. Therefore, when you send an API request to create such VPC and subnet in AWS, NIOS recognizes only the VPC, not the subnet. As a result, you are not able to create VMs under the subnet. For more information about how to create VPCs and subnets in AWS for NIOS, refer to the *Infoblox Installation Guide for vNIOS for AWS*.

In addition, when you delegate authority for supported cloud objects, NIOS may process the requests differently based on the following:

- How the object was first created.
- Whether authority for the object has already been delegated to a Cloud Platform Appliance.

For details about authority delegation and restrictions for each object, see [About Authority Delegation](#).



Note

NIOS does not process cloud API requests that contain unsupported object types or any combination of supported object types with unsupported methods and functions. Although you can use all the fields in a supported object type, some restrictions may apply to supported values for some of these fields. For restrictions, see the **Comments** field in the below table for the corresponding object.

*Supported Cloud API Objects for Cloud API Service*



Note

The cloud API service does not support scheduling and workflow approval requests. Objects deleted through a cloud API request are not stored in the Recycle Bin, except for DNS zones and network views. For information about the Recycle Bin, see [Using the Recycle Bin](#).

Supported Object Type	Cloud API Object	Allowed Operations in cloud API Requests	Authority Delegation and Restrictions	Required Extensible Attributes in cloud API Requests (for creations only)
Network View	networkview	Read, Create, Modify, Delete	See <a href="#">Network Views</a> for information about authority delegation.	Tenant ID Cloud API Owned CMP Type
IPv4 Network Container	networkcontainer	Read, Create, Modify, Delete  Function: next_available_network	Split network, join networks, and RIR related operations are not supported. See <a href="#">IPv4 and IPv6 Networks and Network Containers</a> for information about authority delegation.	Tenant ID Cloud API Owned CMP Type

Supported Object Type	Cloud API Object	Allowed Operations in cloud API Requests	Authority Delegation and Restrictions	Required Extensible Attributes in cloud API Requests (for creations only)
IPv6 Network Container	ipv6networkcontainer	Read, Create, Modify, Delete  Function: next_available_network	Split network, join networks, and RIR related operations are not supported. See <a href="#">IPv4 and IPv6 Networks and Network Containers</a> for information about authority delegation.	Tenant ID Cloud API Owned CMP Type
IPv4 Network	network	Read, Create, Modify, Delete  Function: next_available_ip	Split network, join networks, and RIR related operations are not supported. See <a href="#">IPv4 and IPv6 Networks and Network Containers</a> for information about authority delegation.	Tenant ID Cloud API Owned CMP Type
IPv6 Network	ipv6network	Read, Create, Modify, Delete  Function: next_available_ip	Split network, join networks, and RIR related operations are not supported. See <a href="#">IPv4 and IPv6 Networks and Network Containers</a> for information about authority delegation.	Tenant ID Cloud API Owned CMP Type
IPv4 DHCP Range	range	Read, Create, Modify, Delete  Function: next_available_ip	See <a href="#">DHCP Ranges</a> for information about authority delegation.	Tenant ID Cloud API Owned CMP Type
IPv6 DHCP Range	ipv6range	Read, Create, Modify, Delete  Function: next_available_ip	See <a href="#">DHCP Ranges</a> for information about authority delegation.	Tenant ID Cloud API Owned CMP Type
IPv4 Fixed Address (Reservation)	fixedaddress	Read, Create, Modify, Delete  Function: next_available_ip  You can also create and delete through Grid Manager. All required Cloud EAs are automatically populated in the GUI.	See <a href="#">IPv4 and IPv6 Fixed Addresses</a> for information about authority delegation.	Tenant ID Cloud API Owned CMP Type

Supported Object Type	Cloud API Object	Allowed Operations in cloud API Requests	Authority Delegation and Restrictions	Required Extensible Attributes in cloud API Requests (for creations only)
IPv6 Fixed Address (Reservation)	ipv6fixedaddress	Read, Create, Modify, Delete  Function: next_available_ip  You can also create and delete through Grid Manager. All required Cloud EAs are automatically populated in the GUI.	See <a href="#">IPv4 and IPv6 Fixed Addresses</a> for information about authority delegation.	Tenant ID Cloud API Owned CMP Type
DNS View	view	Read, Modify	See <a href="#">DNS Views</a> for information about authority delegation.	Tenant ID Cloud API Owned CMP Type
DNS Zone	zone_auth	Read, Create, Modify, Delete	See <a href="#">DNS Zones</a> for information about authority delegation.	Tenant ID Cloud API Owned CMP Type
Host Record	record:host	Read, Create, Modify, Delete  You can also create and delete through Grid Manager. All required Cloud EAs are automatically populated in the GUI.	See <a href="#">Host Records</a> for information about authority delegation.	Tenant ID Cloud API Owned CMP Type
	record:host_ipv4addr	Read, Create, Modify, Delete  Function: next_available_ip  You can also create and delete through Grid Manager. All required Cloud EAs are automatically populated in the GUI.		
	record:host_ipv6addr	Read, Create, Modify, Delete  Function: next_available_ip  You can also create and delete through Grid Manager. All required Cloud EAs are automatically populated in the GUI.		



Supported Object Type	Cloud API Object	Allowed Operations in cloud API Requests	Authority Delegation and Restrictions	Required Extensible Attributes in cloud API Requests (for creations only)
Resource Record	record:a	Read, Create, Modify, Delete Function: next_available_ip	See <a href="#">DNS Resource Records</a> for information about authority delegation.	Tenant ID Cloud API Owned CMP Type
	record:aaaa	Read, Create, Modify, Delete Function: next_available_ip		
	record:cname	Read, Create, Modify, Delete		
	record:ptr	Read, Create, Modify, Delete Function: next_available_ip		
	record:mx	Read, Create, Modify, Delete		
	record:naptr	Read, Create, Modify, Delete		
	record:svr	Read, Create, Modify, Delete		
	record:txt	Read, Create, Modify, Delete		
Grid Member	member	Read only Function: restartservices	API requests calling for service restarts on a Grid member can be processed by the Cloud Platform Appliance only if the member requested is also the Cloud Platform Appliance processing the request.	N/A
Grid	grid	Read only Function: restartservices	All cloud API requests calling for service restarts are proxied to the Grid Master.	N/A
Extensible Attribute	extensibleattributedef	Read only	You can use cloud attributes as source objects to obtain the next available IP address or network. When doing so, you must also include the respective network view for the object.	N/A

## Proxying Cloud API Requests

In Cloud Network Automation, the primary Cloud Platform Appliance that receives cloud API requests can act as a proxy for other authoritative Cloud Platform members and for the Grid Master. This proxying mechanism is important when the Cloud Platform Appliance cannot process requests that contain objects for which it is not authoritative, or when objects in the requests do not have authority delegation and must be processed by the Grid Master.

Note that only successfully authenticated and authorized requests that require proxying are sent to the respective appliance for processing. Proxying is limited to one hop within the Grid. Therefore, if the destination appliance cannot process a proxied request, the request will not be forwarded and an error is returned to the client.



### Note

Only cloud API requests can be proxied.

To ensure that the proxying mechanism functions properly, configure your systems to allow for the following communication:

- Allow all HTTPS connectivity among the Cloud Platform Appliances as well as to the Grid Master based on your organization's firewall requirements.
- Ensure that you use the VIP or the MGMT address if it is enabled (including that for the Grid Master) as the destination IP for the HTTPS connectivity. Note that this is a per member setting.
- Grant appropriate permissions to admin groups with cloud API access to ensure that tasks for objects outside of the delegation function properly on the Grid Master.

## Sample Cloud API Requests

This section includes sample cloud API requests for supported objects:

Adding a network view:

```
curl -H "Content-Type: application/json" -k1 -u cloud:infoblox -X POST
https://10.0.0.2/wapi/v2.0/networkview -d '{"name": "netview1", "extattrs":
{ "Tenant ID":{"value": "1011"} , "Cloud API Owned":{"value":"True"},"CMP
Type":{"value":"vCO/vCAC"}}}'
```

Adding a network within the delegated network view in the above example:

```
curl -H "Content-Type: application/json" -k1 -u cloud:infoblox -X POST
https://10.0.0.2/wapi/v2.0/network -d '{ "network":
"20.0.0.0/24", "network_view": "netview1", "extattrs": { "Tenant ID":{"value":
"1011"} , "Cloud API Owned":{"value":"True"},"CMP Type":{"value":"vCO/
vCAC"}}}'
```

Adding a DHCP range within the network created in the above example:

```
curl -H "Content-Type: application/json" -k1 -u cloud:infoblox -X POST
https://10.0.0.2/wapi/v2.0/range -d '{ "end_addr": "20.0.0.40", "member":
{"_struct": "dhcpmember1", "ipv4addr": "10.0.0.2", "name":
" corpxyz.com "}, "network": "20.0.0.0/24", "network_view": "netview1",
"start_addr": "20.0.0.35", "extattrs": {"Tenant ID":{"value": "1011"} , "CMP
Type":{"value":"vCO/vCAC"},"Cloud API Owned":{"value":"True"}}}'
```

Adding an A Record:

```
curl -H "Content-Type: application/json" -k1 -u cloud:infoblox -X POST
https://10.0.0.2/wapi/v2.0/record:a -d '{"name": "corp200.com",
"ipv4addr": "20.0.0.2", "view": "default.netview1", "extattrs": {"Tenant ID":
{"value": "1011"} , "CMP Type": {"value": "vCO/vCAC"}, "Cloud API Owned":
{"value": "True"}, "VM ID": {"value": "12"}}}'
```

Adding a Fixed address:

```
curl -H "Content-Type: application/json" -k1 -u cloud:infoblox -X POST
https://10.0.0.2/wapi/v2.0/fixedaddress -d '{"ipv4addr": "20.0.0.5",
"network_view": "netview1", "mac": "15:06:32:16:00:00", "extattrs": { "Tenant
ID": {"value": "1011"} , "CMP Type": {"value": "vCO/vCAC"} , "VM ID":
{"value": "352"}, "Cloud API Owned": {"value": "True"}}}'
```

Adding a zone:

```
curl -H "Content-Type: application/json" -k1 -u cloud:infoblox -X POST
https://10.0.0.2/wapi/v2.0/zone_auth -d '{ "fqdn": "test.com", "grid_primary":
[{"name": "infoblox.localdomain", "stealth": false}, {"name": "corpxyz.com",
"stealth": false}], "view": "default.netview1", "extattrs": { "Tenant ID":
{"value": "1011"} , "CMP Type": {"value": "vCO/cCAC"} , "Cloud API Owned":
{"value": "True"}}}'
```

Adding a network container:

```
curl -H "Content-Type: application/json" -k1 -u cloud:infoblox -X POST
https://10.0.0.2/wapi/v2.0/networkcontainer -d '{ "network":
"200.0.0.0/24", "network_view": "netview1", "extattrs": { "Tenant ID":
{"value": "1011"} , "Cloud API Owned": {"value": "True"}, "CMP Type":
{"value": "vCO/vCAC"}}}'
```

Add a host record:

```
curl -H "Content-Type: application/json" -k1 -u cloud:infoblox -X POST
https://10.0.0.2/wapi/v2.0/record:host -d '{ "ipv4addrs":
[{"configure_for_dhcp": false, "ipv4addr": "20.0.0.1", "mac":
"11:11:22:22:33:33"}], "ipv6addrs": [{"configure_for_dhcp": false, "duid":
"11:22", "ipv6addr": "13::1"}, {"configure_for_dhcp": false, "duid": "21:22",
"ipv6addr": "13::2"}], "name": "host.corpxyz.com", "view": "default.netview1"}'
```

Adding an MX Record:

```
curl -H "Content-Type: application/json" -k1 -u cloud:infoblox -X POST
https://10.0.0.2/wapi/v2.0/record:mx -d '{ "mail_exchanger":
```

```
" abc.com ", "name": " def.corpxyz.com ", "preference":  
10, "view": "default.netview1", "extattrs": { "Tenant ID": {"value": "1011"} ,  
"CMP Type": {"value": "vCO/vCAC"}, "Cloud API Owned": {"value": "False"}, "VM  
ID": {"value": "230"}}}
```

## Sample Cloud API Requests for Elastic Scaling

### Creating a Member:

```
curl -H "Content-Type: application/json" -k1 -u cloud:cloud -X POST https://  
10.40.240.88/wapi/v2.2/member -d '{"platform": "VNIOS", "host_name":  
"test1.com", "vip_setting": {"address": "1.1.1.1", "gateway": "1.1.0.2",  
"subnet_mask": "255.255.0.0"}}'
```

### Getting a Member:

```
curl -H "Content-Type: application/json" -k1 -u cloud:cloud -X GET https://  
10.40.240.88/wapi/v2.2/member
```

### Adding Pre-Provisioned Information for the Member:

```
curl -H "Content-Type: application/json" -k1 -u cloud:cloud -X PUT https://  
10.40.240.101/wapi/v2.2/member/b25lLnZpcnR1YWxfbm9kZSQ3:test1.com -d  
'{"pre_provisioning": {"hardware_info": [{"hwmodel": "CP-V1405", "hwtype":  
"IB-VNIOS"}]}, "licenses": ["cloud_api", "dhcp", "dns", "enterprise",  
"vnios"]}'
```

### Creating and Delegating a Network View:

```
curl -H "Content-Type: application/json" -k1 -u cloud:cloud -X POST https://  
10.40.240.88/wapi/v2.2/networkview -d '{"name": "testnv", "extattrs": {"Tenant  
ID": {"value": "1011"} , "CMP Type": {"value": "vm130ctest"}, "Cloud API  
Owned": {"value": "True"}} }'
```

### Creating and Delegating a Network:

```
curl -H "Content-Type: application/json" -k1 -u cloud:cloud -X POST https://  
10.40.240.88/wapi/v2.2/network -d '{"network": "21.0.0.0/8",  
"network_view": "default", "cloud_info": {"delegated_member": {"ipv4addr":  
"1.1.1.1", "name": "test1.com"}}, "extattrs": {"Tenant ID": {"value": "1011"} ,  
"CMP Type": {"value": "vm130ctest"}, "Cloud API Owned": {"value": "True"}} }'
```

### Undelegating a Network:

```
curl -H "Content-Type: application/json" -k1 -u cloud:cloud -X PUT https://10.40.240.88/wapi/v2.2/network/ZG5zLm5ldHdvcmskMjEuMC4wLjAvOC8w:21.0.0.0/8/default -d '{"cloud_info": {"delegated_member": null }}'
```

#### Creating and Delegating an Authoritative Zone

```
curl -H "Content-Type: application/json" -k1 -u cloud:cloud -X POST https://10.40.240.88/wapi/v2.2/zone_auth -d '{"fqdn": "test.com", "grid_primary": [{"name": "test1.com", "stealth": false}], "view": "default", "extattrs": {"Tenant ID":{"value": "1011"} , "CMP Type":{"value":"vm130ctest"}, "Cloud API Owned":{"value":"True"}}}'
```

#### Deleting a Member:

```
curl -H "Content-Type: application/json" -k1 -u cloud:cloud -X DELETE https://10.40.240.88/wapi/v2.2/member/b25lLnZpcnR1YWxfbm9kZSQ3:test1.com
```

Creating Token for the Pre-Provisioned Member (note that only superuser can create a token; you must configure superusers admin groups with cloud API access):

```
curl -H "Content-Type: application/json" -k1 -u cloud:cloud -X POST https://10.40.240.88/wapi/v2.2/member/b25lLnZpcnR1YWxfbm9kZSQ3:test1.com?_function=create_token
```

Reading Token for the Pre-Provisioned Member (note that only superuser can create a token; you must configure superusers admin groups with cloud API access):

```
curl -H "Content-Type: application/json" -k1 -u cloud:cloud -X POST https://10.40.240.88/wapi/v2.2/member/b25lLnZpcnR1YWxfbm9kZSQ3:test1.com?_function=read_token
```

## About Authority Delegation

Authority delegation in Cloud Network Automation is the ability to assign full and exclusive control of IP addresses and DNS name spaces to a Cloud Platform Appliance. You can perform authority delegation only through the Grid Master. When you delegate the authority of IP addresses and DNS name spaces to a Cloud Platform Appliance, the Grid Master loses its authority over the scope of delegation for these IP addresses and name spaces as well as any objects within them. Note that authority delegation for an object can be explicitly assigned or inherited from parent objects. For information about how to delegate authority for supported object types, see [Guidelines for Delegating below](#) . NIOS admin users who do not belong to admin groups with cloud API access are not allowed to create new cloud objects, nor can they modify or delete existing cloud objects in delegated spaces; but they can modify the permissions and certain extensible attribute values for these objects. Only admin users with cloud API access and the correct global and object permissions can be used to send cloud API requests to create, modify, and delete objects within the delegated scope.



#### Note

You can delegate authority only to Cloud Platform Appliances, but not to other Grid members.

Objects that are in queue for scheduled executions or approvals are locked and cannot be delegated. Authority delegation and reclaiming of authority are subject to approval and can be scheduled.

## Guidelines for Delegating Authority

You can initiate explicit delegation of authority only through Grid Manager on the Grid Master. The cloud API service can only be used for implicit or automatic delegation of an object such as creating a network under a network container that has been delegated, in which network is implicitly delegated to the member to which the network container is delegated. The Grid Master can explicitly delegate authority only for the following object types:

- Network View
- Network Container (both IPv4 and IPv6)
- Network (both IPv4 and IPv6)
- DHCP Range (IPv4 and IPv6)
- DNS Authoritative Zone (Note that zones are implicitly delegated if the assigned name server is a Cloud Platform Appliance.)

Consider the following when you delegate authority for an object:

- You can delegate authority for supported objects to one and only one Cloud Platform Appliance, except for DNS zones.
- When delegating authority for a parent object, all child objects within the scope of delegation inherit the same authority delegation.
- You cannot delegate authority for the following:
  - Objects whose parents already have a delegation configured.
  - Individual IPAM and DNS records including fixed addresses, host records, A/AAAA/PTR records, etc.
- When you use Elastic Scaling to pre-provision an offline Cloud Platform Appliance, any object authority delegated to this offline member does not take effect until the member joins the Grid. Therefore, you can still create child objects through Grid Manager under the delegated objects when the member is offline.
- You can override the inheritance of authority delegation at the object level only if the parent object has not been delegated. The Grid Master assumes authority for objects that do not fall within the scope of delegation.
- If a supported object has already been delegated, you cannot re-delegate it to another appliance. If you want to re-delegate this object, you must first un-delegate it.
- For explicitly delegated objects, you can only modify the permission and extensible attributes from Grid Manager and the Infoblox API other than cloud API requests. For explicitly delegated zones however, you can modify any properties from Grid Manager and the Infoblox API other than cloud API requests.
- When you create or delete a delegated object through a cloud API request, the appliance returns an OK message if the operation is successful. It returns an ERROR message if the operation fails. You can then change the options in the request and try again. The appliance sends a WARNING message when certain operations require attention.
- You can reclaim the authority that you delegated to a Cloud Platform Appliance. Once the authority is reclaimed, it goes back to the Grid Master. Before you reclaim authority for any object, ensure that the Cloud Platform Appliance is online and properly connected to the Grid Master for the reclaiming process to function properly.
- The Cloud Platform Appliance can run discovery on any network containers or networks that are reachable by the appliance. The default discovery settings for network containers and networks are inherited from their parent objects. For information about discovery, see [About Discovery](#).



### Note

Any Cloud Platform Appliances that are removed from the Grid automatically lose authority over objects that were delegated to them. The Grid Master becomes authoritative for these objects.

## Delegating Authority for Cloud Objects

You can delegate authority when you create a new object that has not been delegated or does not inherit authority delegation from one of its parent objects.

See the following sections for detailed information about delegating authority for supported objects.

## Network Views

Consider the authority delegation guidelines mentioned in the table below when you create, modify, or delete a network view. See [Sample Cloud API Requests](#) for a sample cloud API request.

For information about how to create network views from the Grid Master, see [Adding Network Views](#).

### Authority Delegation for Network Views

Cloud API Requests	Standard API and WAPI Requests	Comments
<ul style="list-style-type: none"> <li>You can delegate authority for a network view to only one Cloud Platform Appliance.</li> <li>When you create a new network view, authority is automatically delegated to the Cloud Platform Appliance that processes the request.</li> <li>To balance network views among multiple Cloud Platform Appliances in the Grid, ensure that you configure your cloud adapter accordingly.</li> <li>If you want to share a network view among different Cloud Platform members, you must manually provision it and its child objects and delegate them to the respective Cloud Platform members.</li> </ul>	<ul style="list-style-type: none"> <li>You can delete a network view from the Grid Master only if it has not been delegated to any Cloud Platform Appliance.</li> <li>When you create a network view on the Grid Master, it is shared among all Grid members in the Grid.</li> <li>You can delegate a network view from the Grid Master to a Cloud Platform Appliance only if the child objects within the network view are delegated to the same Cloud Platform Appliance.</li> <li>When you reclaim authority for a network view, any DNS zones in the network view remain assigned to their name servers, including the Cloud Platform Appliance that has lost authority over the network view. In other words, the DNS zone remains under the authority of that Cloud Platform Appliance.</li> </ul>	<ul style="list-style-type: none"> <li>When you create a network view through a cloud API request, you must include the following extensible attributes in the cloud API request: Tenant ID, Cloud API Owned, and CMP Type.</li> </ul>

## IPv4 and IPv6 Networks and Network Containers

Consider the authority delegation guidelines mentioned in the table below when you create, modify, or delete a network or network container. See [Sample Cloud API Requests](#) for a sample cloud API request. For information about how to create IPv4 and IPv6 networks from the Grid Master, see [Adding IPv4 Networks](#) and [Adding IPv6 Networks](#).

For information about how to create IPv4 and IPv6 networks using network templates from the Grid Master, see [Adding IPv4 Network Templates](#) and [Adding IPv6 Network Templates](#).

### Authority Delegation for Networks and Network Containers

Cloud API Requests	Standard API and WAPI Requests	Comments
<ul style="list-style-type: none"> <li>• You can delegate authority for a network or network container to only one Cloud Platform Appliance, but you can delegate authority of multiple networks and network containers to the same Cloud Platform Appliance.</li> <li>• When you create a new network or network container through a cloud API request, authority is automatically delegated to the primary Cloud Platform Appliance that processes the request.</li> <li>• When you create a network using a network template, you must provide the name of the template and reference it in the cloud API request.</li> <li>• Delegation for a network or network container (except for unmanaged networks) can be done through explicit delegation or inheritance from the parent object.</li> <li>• You cannot delete networks and network containers using cloud API requests if they have already been explicitly delegated. You must first un-delegate them before deleting them.</li> <li>• Networks and network containers associated with a DNS zone cannot be delegated.</li> <li>• You can delegate a network or network container if all the following are true: <ul style="list-style-type: none"> <li>• It has not been delegated to a Cloud Platform Appliance.</li> <li>• It is not part of a network or network container that has been delegated to a Cloud Platform Appliance.</li> <li>• It does not contain any networks or DHCP ranges that are delegated to a different Cloud Platform Appliance.</li> <li>• It does not belong to a delegated network view.</li> </ul> </li> <li>• All discovery related attributes for a network or network container return the default values.</li> <li>• You cannot modify the discovery settings for networks that are in a delegated network container. You cannot create discovered networks in a network container whose authority has been delegated. You also cannot convert unmanaged networks. A discovered unmanaged IP address may co-exist with an IP address created through a cloud API request.</li> <li>• Although no DHCP service restart is required, you can perform a DHCP service restart on a Cloud Platform Appliance through a cloud API request.</li> </ul>	<ul style="list-style-type: none"> <li>• You cannot create, modify, or delete a network or network container on the Grid Master if the object has been delegated to a Cloud Platform Appliance.</li> <li>• You can create a network or network container and delegate it to a Cloud Platform Appliance using a network template.</li> <li>• You can delegate a network or network container to a Cloud Platform Appliance only if the child objects within the network or network container are delegated to the same Cloud Platform Appliance.</li> <li>• You cannot perform a recursive deletion of a network container if one of the child objects in the container has been delegated.</li> <li>• DHCP utilization usage for a network or network container is only updated on the Grid Master.</li> <li>• Discovered IP addresses that are within a delegated network are created on the Grid Master and replicated to the Cloud Platform Appliance that is relevant to the IP addresses.</li> </ul>	<ul style="list-style-type: none"> <li>• When you create a network or network container through a cloud API request, you must include the following extensible attributes in the request: Tenant ID, Cloud API Owned, and CMP Type.</li> <li>• If a network is explicitly delegated (not through inheritance), you can convert a network to a network container only if the size of the network remains the same; the network delegation is transferred to the network container.</li> <li>• The Cloud Platform Appliance does not support split network, john networks, and RIR related operations.</li> </ul>



## DHCP Ranges

Consider the authority delegation guidelines mentioned in the table below when you create, modify, or delete a DHCP range. See [Sample Cloud API Requests](#) for a sample cloud API request.

For information about how to create IPv4 and IPv6 ranges, see [Adding IPv4 Address Ranges](#) and [Modifying IPv6 Address Ranges](#).

For information about how to create IPv4 and IPv6 ranges using range templates, see [Adding IPv4 Range Templates](#) and [Adding IPv6 Range Templates](#).

*Authority Delegation for DHCP Ranges*

Cloud API Requests	Standard API and WAPI Requests	Comments
<ul style="list-style-type: none"> <li>• You can delegate authority for a DHCP range to only one Cloud Platform Appliance, but you can delegate authority for multiple DHCP ranges to the same Cloud Platform Appliance.</li> <li>• When you create a new DHCP range, authority is automatically delegated to the Cloud Platform Appliance that processes the request or to the Grid Master if the Grid Master processes the request.</li> <li>• When you create a DHCP range using a range template, you must know the name of the template and reference it in the cloud API request.</li> <li>• You can delegate authority for reserved ranges in a Microsoft synchronized network if: <ul style="list-style-type: none"> <li>• It is not included in any exclusions.</li> <li>• It does not conflict with another reserved range.</li> </ul> <p>You can manage these range addresses through a cloud adapter, such as the IPAM Plug-In for VMware.</p> </li> <li>• Delegation for a DHCP range can be done through explicit delegation or inheritance from the parent object. You cannot override inherited delegation.</li> <li>• Note that you cannot delete DHCP ranges using cloud API requests if they have already been explicitly delegated. You must first un-delegate them before deleting them. However, if the delegation is inherited, you can delete the objects through a cloud API request.</li> <li>• You can delegate a DHCP range to a Cloud Platform Appliance if all the following are true: <ul style="list-style-type: none"> <li>• It has not been delegated to a Cloud Platform Appliance.</li> <li>• It is not part of a delegated network or network container in the same network view.</li> <li>• It does not belong to a delegated network view.</li> <li>• It is a reserved range or a range that has been assigned to a DHCP member that is the same Cloud Platform Appliance to which you want to delegate the range.</li> </ul> </li> <li>• You cannot delegate a DHCP range that has been assigned to a failover association, nor can you assign a DHCP range that has been delegated to a failover association.</li> <li>• Authority is delegated from the start address to the end address, including exclusions. Note that the exclusions can be used only to restrict IP addresses generated by the next available IP feature.</li> <li>• All discovery related attributes for a DHCP range return the default values.</li> <li>• Although no DHCP service restart is required, you can perform a DHCP service restart on a Cloud Platform Appliance through a cloud API request.</li> </ul>	<ul style="list-style-type: none"> <li>• You cannot create, modify, or delete a DHCP range on the Grid Master if it has been delegated to a Cloud Platform Appliance.</li> <li>• You can create a DHCP range and delegate it to a Cloud Platform Appliance using a range template.</li> <li>• You can increase the size of a DHCP range that has been explicitly delegated to a Cloud Platform Appliance. The increased size is available for use by the Cloud Platform Appliance after replication.</li> </ul>	<ul style="list-style-type: none"> <li>• When you create a DHCP range through a cloud API request, you must include the following extensible attributes in the request: Tenant ID, Cloud API Owned, and CMP Type.</li> </ul>

Cloud API Requests	Standard API and WAPI Requests	Comments

## IPv4 and IPv6 Fixed Addresses

Consider the following authority delegation guidelines when you create, modify, or delete a fixed address:

- You can delegate authority for a fixed address only through inheritance from one of its parent objects, such as its associated network view, network container, network, or DHCP reserved range.
- When you create or modify an IPv4 or IPv6 fixed address, you must include the following extensible attributes in the cloud API request: Tenant ID, Cloud API Owned, and CMP Type.
- You can create a fixed address from the Grid Master using a fixed address template. Note that when you want to reference a template in the cloud API request, you must know the name of the template beforehand.
- When performing any operations on a Cloud Platform Appliance, all discovery related attributes for a fixed address return the default values.
- No DHCP service restart is required when performing any operations for a fixed address on the Cloud Platform Appliance unless automatic DHCP restart is disabled on the appliance. You can however perform a DHCP service restart on the Cloud Platform Appliance to which authority is delegated for a fixed address through a cloud API request.
- You can create, modify, or delete an IPv4 or IPv6 fixed address and reservation on the Grid Master through Grid Manager if the fixed address or reservation is within the scope of a network view, network container, network, or DHCP reserved range whose authority has been delegated to a Cloud Platform Appliance.

See [Sample Cloud API Requests](#) for a sample cloud API request.

For information about how to create IPv4 and IPv6 fixed addresses, see [Adding IPv4 Fixed Addresses](#) and [Adding IPv6 Fixed Addresses](#).

For information about how to create IPv4 and IPv6 fixed address templates, see [Adding IPv4 Fixed Address/Reservation Templates](#) and [Adding IPv6 Fixed Address Templates](#).

## DNS Views

Consider the following authority delegation guidelines when you create, modify, or delete a DNS view:

- You cannot explicitly delegate authority for a DNS view. The Cloud Platform Appliance automatically gains authority over any DNS view that exists in the network view whose authority is delegated to that appliance.
- You cannot create or delete a DNS view from the Cloud Platform Appliance.
- Through a cloud API request, you can update DNS views defined in a network view that has been delegated to the Cloud Platform Appliance.
- You cannot create, modify, or delete a DNS view in network views that have been delegated to a Cloud Platform Appliance through a standard API request.
- You cannot delete a DNS view as long as it contains at least one DNS zone that has been delegated to a Cloud Platform Appliance.

## DNS Zones

Consider the following authority delegation guidelines mentioned in the table below when you create, modify, or delete a DNS zone. See [Sample Cloud API Requests](#) for a sample cloud API request.

For information about how to create DNS zones, see [Configuring Authoritative Zones](#).

### *Authority Delegation for DNS Zones*

Cloud API Requests	Standard API and WAPI Requests	Comments
<ul style="list-style-type: none"> <li>The Grid primary of a DNS zone automatically gains authority for the zone if the primary is a Cloud Platform Appliance. When there are multiple primaries configured for the zone, multiple delegations to these primaries are allowed as long as they are Cloud Platform Appliances.</li> <li>You cannot assign both a Microsoft server and a Grid member as primaries at the same time, although you can assign a Microsoft server as the Grid primary and a Cloud Platform Appliance as the Grid secondary. This allows the Microsoft server to serve changes sent from the cloud adapter.</li> <li>All resource records in a DNS zone inherit authority delegation from the zone. However, you cannot modify the NS record through a cloud API request.</li> <li>You can modify all the fields for a zone whose authority has been explicitly delegated.</li> <li>The cloud member to which authority for a network view is delegated automatically gains authority for authoritative zones defined in that network view. This Cloud Platform Appliance is the only cloud member that can be the Grid primary for the zones defined in this network view. The Grid Master does not have authority for any zone in this network view unless it is assigned as a Grid primary.</li> <li>The Cloud Platform Appliance can create, modify, and delete a DNS zone in any DNS view defined in a network view whose authority has been delegated to that cloud member.</li> <li>The Cloud Platform Appliance that is authoritative for a DNS zone can perform changes to the assigned Grid primaries, Grid secondaries, and external servers assigned to the zone as long as the Cloud Platform Appliance remains a Grid primary. But it cannot create, modify, or delete the NS record.</li> <li>The Cloud Platform Appliance that is authoritative for a DNS zone can create, modify, and delete DNS delegations that are directly parented to that zone. In particular, it may specify any Grid primary, Grid secondary, or external server for that zone.</li> <li>DNSSEC operations, network associations, and zone locking are not supported if at least one Cloud Platform Appliance is assigned as the Grid primary for any DNS zones.</li> <li>Although no DHCP service restart is required, you can perform a DHCP service restart on a Cloud Platform Appliance through a cloud API request.</li> </ul>	<ul style="list-style-type: none"> <li>You cannot create, modify, or delete a DNS zone in a network view whose authority has been delegated to a Cloud Platform Appliance.</li> <li>You cannot assign a Cloud Platform Appliance as the Grid primary for a zone that is locked or disabled.</li> <li>You can modify extensible attributes of any DNS zone whose authority has been delegated from the Grid Master.</li> </ul>	<ul style="list-style-type: none"> <li>Only authority for authoritative forward-mapping and reverse-mapping zones can be delegated. You cannot delegate authority for forward zones, stub zones, and delegated zones even though they may exist in a delegated network view.</li> <li>When you create a DNS zone using a cloud API request, you must include the following extensible attributes in the request: Tenant ID, Cloud API Owned, and CMP Type.</li> </ul>

## DNS Resource Records

Consider the following authority delegation guidelines mentioned in the table below when you create, modify, or delete a resource record, including A, AAAA, CNAME, PTR, MX, SRV, TXT, NAPTR records.

See [Sample Cloud API Requests](#) for a sample cloud API request.

### Authority Delegation for DNS Resource Records

Cloud API Requests	Standard API and WAPI Requests	Comments
<ul style="list-style-type: none"><li>• Authority delegation for resource records (A, AAAA, CNAME, PTR, MX, SRV, TXT, and NAPTR) is inherited from their parent zones. You can delegate authority for these records by delegating authority for their respective parent zones.</li><li>• All resource records in a DNS zone inherit authority delegation from their parent zones. However, you cannot modify the NS record through a cloud API request.</li><li>• If the Cloud Platform Appliance is a Grid primary for a zone, requests that include a supported record is processed locally by the Cloud Platform Appliance. Otherwise, the request is proxied to the Cloud Platform Appliance that is assigned as the only Grid primary for the zone.</li><li>• If the DNS resource records belong to a zone that is served only by Cloud Platform Appliances, authority for these records are considered delegated. You must create, modify, or delete these records on one of these Cloud Platform Appliances.</li></ul>	<ul style="list-style-type: none"><li>• You cannot create, modify, or delete a resource record if it is in a network view whose authority has been delegated to a Cloud Platform Appliance.</li></ul>	<ul style="list-style-type: none"><li>• When you create a resource record through a cloud API request, you must include the following extensible attributes in the request: Tenant ID, Cloud API Owned, and CMP Type.</li></ul>

### Host Records

Consider the following authority delegation guidelines mentioned in the table below when you create, modify, or delete a host record. See [Sample Cloud API Requests](#) for a sample cloud API request.

#### Authority Delegation for Host Records

Cloud API Requests	Standard API and WAPI Requests	Comments
<ul style="list-style-type: none"> <li>Authority delegation for a host record is inherited from both the DNS and DHCP portions of the record. For DNS, you can delegate authority for all DNS zones for which the host record is defined. For DHCP, you can delegate authority for the parent network view, network container, network, or DHCP range defined for the host record.</li> <li>You can create, modify, or delete a host record or a host IP address whose authority is delegated to a Cloud Platform Appliance through Grid Manager. Note that when you create a host record, you must enable it for DNS within the delegated network view. Otherwise, you will not be able to save the host record.</li> <li>The Cloud Platform Appliance can process a cloud API request that includes a host record only if it has gained authority for both DNS and DHCP portions of the host record, as follows: <ul style="list-style-type: none"> <li>All IP addresses enabled for DHCP within one or more delegation scopes are delegated to the same Cloud Platform Appliance.</li> <li>All DNS records defined for one or more DNS zones have the same Cloud Platform Appliance assigned as the Grid primary.</li> </ul> </li> </ul>	<ul style="list-style-type: none"> <li>IP addresses defined in the host record that is enabled DHCP follow the same rules for a fixed address. See <a href="#">IPv4 and IPv6 Fixed Addresses</a> for more information.</li> <li>Names or aliases defined in the host record follow the same rules set for resource records. See <a href="#">DNS Resource Records</a> for more information.</li> </ul>	<ul style="list-style-type: none"> <li>When you create a host record through a cloud API request, you must include the following extensible attributes in the request: Tenant ID, Cloud API Owned, and CMP Type.</li> </ul>
<ul style="list-style-type: none"> <li>IP addresses defined in the host record that is enabled for DHCP follow the same rules set for a fixed address. See <a href="#">IPv4 and IPv6 Fixed Addresses</a> for more information.</li> <li>Names or aliases defined in the host record follow the same rules set for resource records. See <a href="#">DNS Resource Records</a> for more information.</li> <li>Although no DHCP service restart is required, you can perform a DHCP service restart on a Cloud Platform Appliance through a cloud API request.</li> </ul>		

## Configuring Grid and Member Cloud API Properties

Only admin users in admin groups with cloud API access can be used to send cloud API queries by default. For information about how to add users to an admin group, see [Creating Local Admins](#). To control which admin users,

either from this group or from remote servers, can perform cloud API tasks, you can further define ACLs at the Grid and member levels.

1. **Grid:** From the **Cloud** tab, click **Grid Cloud API Properties** from the Toolbar. Configuration done for the Grid only applies to the current Grid Master; it is not inherited by other Cloud Platform Appliances.

**Member:** From the **Cloud** tab, select the **Members** tab -> *member* checkbox, and then click the Action icon

and select **Edit** from the menu. Configuration done at the member level applies only to the Grid member.

2. In the *Grid Cloud API Properties* (for the current Grid Master) or the *Member Cloud API Properties* editor, select the **General** tab, and then complete the following:

**Administrator allowed to make WAPI request on the Grid Master**

- **None:** When you select this, none of the admin users in admin groups with cloud API access can send cloud API requests to the Grid Master or Cloud Platform Appliance.
- **All:** When you select this, all admin users in admin groups with cloud API access can send cloud API requests to the Grid Master or cloud Platform Appliance. This is the default.
- **Set of administrators:** Select this to create a list of admin users, both remote and local, who can send cloud API requests. Local users are users defined in admin groups with cloud API access. Remote users are users who log in from other remote servers. These users will be authenticated before they can access the Grid Master or Cloud Platform Appliance.

To add local users, click the Add icon and select **Local**. In the *Cloud API Admin Selector*, select an admin user from the list. Grid Manager adds the selected user to the table. If you have only one cloud API user, Grid Manager automatically adds this user to the table.

To add remote users, click the Add icon and select **Remote**. Grid Manager adds a row to the table. Click the **Admin** column to add the username of the administrator. Note that the username you enter here must match the username used on the remote server. Depending on the remote server type, you must create a server group for these remote users and add the group to the admin authentication policy to ensure these admin users can send cloud API requests. For information about how to configure admin server groups and admin authentication policy, see [About Remote Admins](#).

Click the Add icon again to add additional admin users.

**Recycle cloud objects:** This only applies to the Grid Master. Select this checkbox to enable the recycling of cloud objects. This is selected by default.

3. Save the configuration.

## Extensible Attributes for Cloud Objects

When you first enable Cloud Network Automation, NIOS installs a set of extensible attributes that are specific for cloud usage. Use these cloud extensible attributes to tag objects that belong to the CMP. Note the following when defining cloud extensible attributes through Grid Manager:

- You cannot assign cloud extensible attributes to other NIOS objects, yet you can create smart folders using these cloud attributes or modify their definitions.
- You can define and update cloud extensible attributes on Grid Master, through cloud API requests or Grid Manager, as long as the authority for the corresponding cloud objects are not delegated.
- Existing extensible attributes are automatically enrolled for cloud usage when cloud licensed are installed.
- All cloud extensible attributes are displayed in the **Administration** tab -> **Extensible Attributes** tab in Grid Manager.

To identify a cloud object, you must reference some of these cloud attributes when you create, modify, or delete a specific object. For more information about which extensible attributes are required for cloud API requests, see [Supported Cloud API Objects](#).

The following table lists the default cloud extensible attributes come installed on the appliance. Note that some of the attributes are read-only and you cannot modify their properties. These attributes are applicable for specific object types or for identification purposes. See **Comments** for more information about each attribute.

### *Extensible Attributes for Cloud Usage*

Attribute Name	Attribute Type	Comments
Account	String	The CMP user account for creating networks.

Attribute Name	Attribute Type	Comments
Allocation ID	String	The allocation ID of the Elastic IP. Limited to Elastic IP only. Example: eipalloc-5723d13e.
Application Type	String	Indicates the application type, such as Web, DB, or CRM.
Association ID	String	Association ID specific to Elastic IP only.
Attachment ID	String	The attachment ID of the network interface. This is valid for Elastic IP only and present when Network Interface is attached to an instance. Example: eni-attach-d94b09b0.
Availability Zone	String	
Cloud API Owned	List [True, False]	This is read-only. Defines whether an object was created by the cloud adapter.
Cloud Region	String	A region name for an VPC object. Example: us-west-1.
CMP Type	String	This is read-only. Defines the type of CMP, such as VMware or OpenStack.
Host Aggregates	String	
Interface Name	String	The name of the interface.
Is External	List [True, False]	This is read-only. Limited to the object type Network and Network Container.
Is Primary Interface	List [True, False]	This is read-only.
Is Shared	List [True, False]	This is read-only. Limited to the object type Network and Network Container.
IP Type	List [Private, Public, Fixed, Floating, Elastic]	This is read-only. Type of IP address.
Location	String	
Network Encap	String	
Network ID	String	Network ID in OpenStack
Network Name	String	Network name
Physical Network Name	String	
Port Attached Device - Device ID	String	Device ID for associated device, such as OpenStack or equivalent, in other CMPs.



Attribute Name	Attribute Type	Comments
Port Attached Device - Device Owner	String	Device name for associated device, such as OpenStack or equivalent, in other CMPs (e.g. compute:nova, network:dhcp, or network:router_interface).
Port Group	String	VMware or equivalent in other Hypervisors or CMPs.
Port ID	String	Port ID for associated device, such as OpenStack or equivalent, in other CMPs.
Port Name	String	Port name for associated device, such as OpenStack or equivalent, in other CMPs.
Private IP	String	One or more secondary private IP addresses that are assigned to the given Network Interface.
Segmentation ID	String	
Subnet ID	String	
Subnet Name	String	
Tenant ID	String	This is read-only. The unique ID for the tenant object.
vDC	String	
VLAN ID	Integer	The VLAN ID.
VM ID	String	This is read-only. This is the Instance ID in OpenStack.
VM Name	String	Instance Name in OpenStack.
vCD App	String	The application name defined in vCloud Director; previously vApp
vCD Org	String	The organization name defined in vCloud Director; previously vOrg.
VPC ID	String	This is read-only. This is a naming convention that is used at the time of creation. Example: vpc-1a2b3c4d.
VPC Name	String	An optional name tag for the VPC.
VPCs List	String	List of all the VPCs.

You can modify some of the properties for the cloud extensible attributes, except for the read-only attributes. By default, all cloud extensible attributes are configured to allow Read/Write access for the Cloud Platform Appliances. You can change this configuration to read-only so the Cloud Platform Appliances can only access the attribute values, but not modify them. Note that when you reference modification for a read-only attribute in a cloud API request, the Cloud Platform Appliance returns an error because it cannot modify the attribute value. For information about how to configure extensible attributes, see [About Extensible Attributes](#).



#### Note

An upgrade could fail if the name of an existing extensible attribute matches the name of any of the cloud extensible attribute for a different object type. You must define values for all required cloud extensible attributes in a cloud API request.

## Extensible Attributes for Tags in AWS and Azure

You can define metadata in the form of tags for AWS and Azure which are captured through a vDiscovery process and you can save the tags as extensible attributes in NIOS. You can use predefined attributes or create your own tags that consists of an user-defined key and an optional value. The tag values defined in AWS and Azure are translated into corresponding extensible attribute values in NIOS.

Note the following about saving tags defined in AWS and Azure as extensible attributes in NIOS:

- You must add extensible attributes in NIOS with the same name as the tags added in AWS or Azure.
- You can add or delete tags in Azure, but cannot update the tags.
- The tags are translated only when the corresponding extensible attributes are created in NIOS.
- The tags are discovered only during the next vDiscovery process after creating the corresponding extensible attributes in NIOS.
- It is not recommended to delete the extensible attributes which you have created for the tags defined in AWS and Azure.

The following table shows the translation of tags defined in AWS and Azure, as extensible attributes in NIOS:

### Tag Translation

AWS Object	Azure Object	NIOS Object
EC-2 Instance	Virtual Machine	VM
Interface	Virtual Interface	Managed private IP address: Any DNS record, fixed address, or reservation associated with that IP address.
Interface (tags are the same for private IP address and public IP address of the same interface)	Public IP address (Public IP address has specific tags in Azure)	Managed public IP address: Any DNS record, fixed address, or reservation associated with the IP address.
VPC	Virtual Network	VPC
Subnet	Subnet (no tags for subnet in Azure)	Network



#### Note

NIOS generates alert messages about tags that are translated and tags that are skipped due to missed extensible attributes or incorrect extensible attributes types will be displayed in the syslog and infoblox.log file.

## Viewing Cloud Objects

When you enable the Cloud Network Automation license on the Grid Master, NIOS adds the ability to view new cloud objects such as Tenants and VM IP addresses. You can view cloud objects and their related information in the **Cloud** tab of Grid Manager. The **Cloud** tab provides the following sub tabs for viewing different information related to cloud objects: **Tenants, VPCs, Networks, VMs (by IP Address), and Cloud Platform Members.**

In addition to viewing data in these tabs, you can do the following:

1. Click the Action icon  
and select an action from the menu.
  - Select **Show Active Users** to view all the users who are currently active on the Active Directory domain.  
For information, see [Viewing Active Network Users](#).
2. Modify some of the data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes. Note that some fields are read-only.
3. Edit the properties of an object.
  - Select the object, and then click the Edit icon.
4. Export the list of objects to a .csv file.
  - Click the Export icon.
5. Print the list of objects.
  - Click the Print icon.
6. Use filters and the **Goto** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Goto** field and select the object from the possible matches.
7. Create a quick filter to save frequently used filter criteria:
  - In the filter section, click Show Filter and define filter criteria for the quick filter.
  - Click **Save** and complete the configuration In the *SaveQuickFilter* dialog box.

The appliance adds the quick filter to the quick filter drop-down list in the panel. Note that global filters are prefixed with [G], local filters with [L], and system filters with [S].

## Viewing All Tenants

The **Tenants** tab lists all tenants from the CMP. Tenant is an abstract administrative concept. Similar to a tenant in the CMP, a tenant object in NIOS encompasses all network elements such as networks, zones, VMs, and IP addresses (fixed and floating), network views, default DNS view, and all related extensible attributes associated with that tenant. Multiple tenants can be mapped to the same network view. A tenant can also have VMs (with IP addresses) in its shared or external networks that are in different network views.

In Grid Manager, you can click a tenant name in the **Cloud** tab -> **Tenants** tab and drill down to the **Networks** and **VMs (by IP Address)** sub tabs to view networks and VMs associated with the selected tenant. In the **Tenants** tab -> **VMs** tab, you can click a VM name and drill down more to view the **Networks** and **IP Addresses** sub tabs for the selected VM. You can always click the bread crumb at the top of the viewer to go back to the **Tenants** home tab.

Each tenant has a name and a unique tenant ID. The tenant ID is provided through cloud API requests. You cannot create or delete tenant objects through Grid Manager. All tenants are created and deleted through cloud API requests. However, you can modify the name, extensible attributes, and permissions for a specified tenant through the *Tenant* editor in Grid Manager if you have valid tenant permissions: **All Tenants** or per tenant object. Note that you cannot delegate the authority of any given tenants.

To view all tenant objects:

1. From the **Cloud** tab, click the **Tenants** tab.
2. Grid Manager displays the following information for each tenant:
  - **Actions:** Click the action icon  
(shown as a gear in each row) next to a selected tenant and choose from the following:
    - **Edit:** Modify certain general properties.
    - **Extensible Attributes:** Add or modify extensible attributes.
    - **Permissions:** Modify the administrative permissions.
  - **Mgmt Platform:** Displays the CMP that manages this tenant. When it displays **Amazon**, it indicates a successful validation of the Amazon account from NIOS to AWS.

- **Name:** The tenant name.
- **ID:** The unique tenant ID.
- **VMs:** The total number of VM objects associated with this tenant. This can include the following object types: Host Record, Fixed Address, and any resource record type such as A, AAAA, PTR, and CNAME records. It also includes unmanaged IP addresses that are associated with the tenant.
- **Networks:** The total number of IPv4 and IPv6 networks and network containers associated with this tenant. Note that this number includes only networks and network containers created by the cloud adapter.
- **Created:** The timestamp when the tenant was first created. You cannot modify this field. This timestamp reflects the time when the tenant object was first seen by the Grid Master, so it may not match the timestamp when the original cloud API request was sent.
- **Last Updated:** The timestamp when the last event associated with this tenant happened. You cannot modify this field. This timestamp reflects the time when the last event associated with this tenant was processed by the Cloud Platform Appliance, so it may not match the timestamp when the original cloud API request was sent.
- **Comment:** Information about this tenant.
- **Network Views:** The network view to which this tenant belongs.
- **Managed:** Indicated whether this tenant is a managed or an unmanaged object in NIOS.
- **Site:** The value entered for this predefined extensible attribute.

You can also select other cloud extensible attributes for display by clicking the down arrow next to any column header and selecting **Columns** -> **Edit Columns**.



#### Note

The vDiscovery for the OpenStack management platform discovers all tenants if the OpenStack user has the admin role in at least one tenant.

## Viewing All VPCs (Virtual Private Clouds)

The **VPCs** tab displays all AWS VPCs. You can also manage selected VPCs, primarily for changing permissions, defining or changing extensible attributes, and changing the delegation settings for a VPC to a different NIOS Cloud member. An Amazon VPC is analogous to a network container in NIOS, and is hence represented as a network container with a special icon.

To view all VPCs:

1. From the **Cloud** tab, click the **VPCs** tab.
2. Grid Manager displays the following information for each VPC:
  - **Actions:** Click the action icon  
  
(shown as a gear in each row) next to a selected tenant and choose from the following:
    - **Edit:** Modify certain general properties.
    - **Extensible Attributes:** Add or modify extensible attributes.
    - **Permissions:** Modify the administrative permissions.
  - **Mgmt Platform:** Displays the CMP that manages the VPC. When it displays **Amazon**, it indicates a successful validation of the Amazon account from NIOS to AWS.
  - **VPC Name:** The AWS virtual private cloud name. The name is automatically defined by AWS. Each VPC name is a link that opens the Networks tab for the selected VPC. This page lists the individual private networks that exist within the VPC.
  - **Networks:** The number of individual private networks contained in the VPC.
  - **VMs:** The number of Amazon EC2 virtual machine instances currently discovered in the VPC. (You can run a vDiscovery in any VPC.) For information about how to start a vDiscovery, see [Configuring vDiscovery Jobs](#).
  - **Tenants:** The number of cloud tenants associated with each VPC.
  - **Cloud Usage:** indicates whether the VPC is associated with any specific cloud extensible attributes or within a scope of delegation. It can be one of the following:
    - **Cloud from adapter:** Indicates that this object has been created by a cloud adapter and it may or may not be within a scope of delegation at the moment.

- **Cloud from delegation:** Indicates that this object is within the scope of delegation or the object itself defines a scope of authority delegation, and it is not created by a cloud adapter.
- **Used by cloud:** Indicates that this network or network container is associated with the extensible attribute **Is External** or **Is Shared** and the value is set to True, which implies the network is a private or shared network managed by the CMP, and it is not **Cloud from adapter** or **Cloud from delegation**.
- **Non-cloud:** The object is a regular NIOS object and is not within the scope of any authority delegation nor is it associated with any of these extensible attributes: **Cloud API Owned**, **Is External** or **Is Shared**. NIOS admin users can modify this object based on their permissions.
- **Owned By:** A cloud object can be owned by the Grid Master or the cloud adapter. When the object is created by the Grid Master, this shows Grid. If the object is created by the cloud adapter, this shows Cloud Adapter.
- **Delegated to:** The NIOS Cloud appliance to which management of the AWS VPC is delegated. This field tells you whether or not a cloud object (in this case, a virtual private cloud) has been delegated to a Cloud Platform Appliance.
- **Network:** The network IP. The network listed in this column for the VPC is also viewable from the main **Data Management** → **IPAM** tab.
- **Site:** Extensible Attribute listing the site information for the VPC.
- **Availability Zone:** the Amazon availability zone in which the VPC resides.

You can also select other cloud extensible attributes for display by clicking the down arrow next to any column header and selecting **Columns** -> **Edit Columns**.

## Viewing All Cloud Networks

The **Networks** tab displays all IPv4 and IPv6 networks and network containers from the CMP as well as all delegated networks. You can select a specific network or network container and modify its properties in the *Cloud IPv4 Network* or *Cloud IPv6 Network* editor.

To view all cloud networks and network containers:

1. From the **Cloud** tab, click the **Networks** tab.
2. Grid Manager displays the following information for each network and network container:
  - **Actions:** Click the action icon  
(shown as a gear in each row) next to a selected tenant and choose from the following:
    - **Go to Tenant:** Go to the **Tenant** tab to view associated tenant.
    - **Go To DHCP Network Details:** Go to the **DHCP** -> **Networks** tab to view associated details.
    - **Go To IPAM Network Details:** Go to the **IPAM** -> **Networks** tab to view associated details.
    - **Go To Network View Details:** Go to the **IPAM** -> **Network View** tab to view associated details.
    - **Edit:** Modify certain general properties.
    - **Extensible Attributes:** Add or modify extensible attributes.
    - **Permissions:** Modify the administrative permissions.
  - **Mgmt Platform:** Displays the CMP that manages the network. When it displays Amazon, it indicates a successful validation of the Amazon account from NIOS to AWS.
  - **Network:** The IP address and netmask of the network.
  - **Tenant:** The associated tenant for the network.
  - **VPC Name:** The name of the associated VPC in AWS.
  - **Cloud Usage:** This field indicates whether this object is associated with any specific cloud extensible attributes or within a scope of delegation. It can be one of the following:
    - **Cloud from adapter:** Indicates that this object has been created by a cloud adapter and it may or may not be within a scope of delegation at the moment.
    - **Cloud from delegation:** Indicates that this object is within the scope of delegation or the object itself defines a scope of authority delegation, and it is not created by a cloud adapter.
    - **Used by cloud:** Indicates that this network or network container is associated with the extensible attribute **Is External** or **Is Shared** and the value is set to True, which implies the network is a private or shared network managed by the CMP, and it is not **Cloud from adapter** or **Cloud from delegation**.

- **Non-cloud:** The object is a regular NIOS object and is not within the scope of any authority delegation nor is it associated with any of these extensible attributes: **Cloud API Owned, Is External** or **Is Shared**. NIOS admin users can modify this object based on their permissions.
- **Owned By:** A cloud object can be owned by the Grid Master or the cloud adapter. When the object is created by the Grid Master, this shows **Grid**. If the object is created by the cloud adapter, this shows **Adapter**.
- **Delegated To:** This tells you whether a cloud object has been delegated to a Cloud Platform Appliance or not. If the cloud object has a parent object and the parent has been delegated, this field shows the parent delegation and you cannot modify the field.
- **Network View:** The network view to which this network belongs.
- **Active Users:** Displays the number of active users on the selected network.
- **Site:** The value entered for this predefined extensible attribute.

You can also select other cloud extensible attributes for display by clicking the down arrow next to any column header and selecting **Columns** -> **Edit Columns**.

## Viewing All Cloud VMs

The **VMs** tab lists all cloud VMs by IP addresses. A VM object is an abstract object representing a virtual machine that is running on the CMP. A VM belongs to one and only one tenant. However, the same VM may have more than one IP addresses (including unmanaged IP addresses) associated with it. Each VM may have an IP address that is part of an overlapping private IP address space and one or more IP addresses in the shared or external networks.

A VM object in NIOS can be a collection of supported object types that have the same value for the extensible attribute **VM ID**. Only the following NIOS object types are considered as existing VMs when they are tagged with the **VM ID** attribute: Host Record, A Record, AAAA Record, and PTR Record, and Fixed Address. A VM object may be defined by objects from different network views, and it can have more than one IP addresses associated with it.

**Note:** Since a VM can be defined by objects from different network views, the same IP address may appear multiple times if it has been defined in more than one network view. A VM object is a read-only abstract object, therefore you cannot create, modify, or delete it.

After a vDiscovery job is completed, the appliance displays discovered data for each VM in this tab. Available data is displayed based on the vDiscovery configuration and your CMP. For example, if your CMP is AWS, discovered data can include the VPC to which the VM belongs. You can click a VM name and drill down to the **Networks** and **IP Addresses** sub tabs to view networks and IP addresses associated with the selected VM. For more information about vDiscovery, see [Configuring vDiscovery Jobs](#).

Note that in addition to managing discovered data through Grid Manager, you can clear any managed or unmanaged discovered data, or clear all discovery data related to a vDiscovery job through a cloud API request. You can use this feature to properly identify VMs that you spin up or de-provision through a cloud adapter. For example, when you use Infoblox IPAM Plug-In for VMware as the cloud adapter to de-provision a VM, you can send a cloud API call to remove the discovered data for this VM so you can avoid IP address conflict with IP addresses manually allocated by the VMware vCenter. For information about cloud API requests, see [About Cloud API Requests](#).

In the **VMs** tab, discovered VMs are highlighted in different background colors, as follows:

- **Yellow:** Unmanaged VMs that do not have associated NIOS objects.
- **White:** Discovered VMs that have at least one associated NIOS object and there is no conflicting information between the discovered data and the NIOS data.
- **Red:** Discovered VMs that have at least one associated NIOS object and there is conflicting information between the discovered data and the NIOS data. Depending on the nature of the conflict, you can resolve them as described in [Resolving Conflicting Addresses](#). You may also be able to convert or clear unmanaged data, as described in [Managing Unmanaged Data](#).

To view all VM objects in NIOS:

1. From the **Cloud** tab, click the **VMs** tab.
2. Grid Manager displays the following information for all cloud VM by IP address:
  - **Actions:** Click the action icon

(shown as a gear in each row) next to a selected tenant and select the action you want to perform.

- **Mgmt Platform:** Displays the CMP to which this tenant belongs. This can be Amazon, OpenStack, or VMware.
- **VM Name:** The name of the VM.
- **VM ID:** The unique tenant ID to which this VM belongs.
- **Networks:** The number of networks that belong to this VM.
- **IP Address:** The IP address of the VM.
- **VM VPC:** The VPC to which this VM belongs.
- **VM Tenant:** The tenant to which this VM belongs.
- **Port ID:** The port ID for the VM.
- **Network View:** The network view to which this VM belongs.
- **Active Users:** The number of active users on the selected network.
- **FQDN:** The FQDN of the VM.
- **VM Last Updated:** The timestamp when the VM data was last updated.
- **First Discovered:** The timestamp when the VM was first discovered.
- **Last Discovered:** The timestamp when the VM was last discovered.
- **Task Name:** The name of the task that collected the discovered data. It is usually the ID or task name that collected the discovered data.
- **Comment:** Information about the VM.

Depending on you CMP, you can also select additional discovered fields to be displayed in the **VMs** tab by clicking the down arrow next to any column header and selecting **Columns -> Edit Columns**. Note that some of these fields contain discovered data that is only relevant to your CMP.

## Viewing All Cloud Platform Members

The **Members** tab displays all members that are currently running the cloud API service. To view all cloud members in NIOS:

1. From the **Cloud** tab, click the **Members** tab.
2. Grid Manager displays the following information for each member:
  - **Actions:** Click the action icon  
(shown as a gear in each row) next to a selected tenant and select the action you want to perform.
  - **Name:** The member name.
  - **Status:** The current status of this member.
  - **Comment:** Information about this cloud member.
  - **Site:** The value entered for this predefined extensible attribute.

Select other cloud extensible attributes for display by clicking the down arrow next to any column header and selecting **Columns -> Edit Columns**.

## Managing Appliance Operations

Managing the operations of a NIOS appliance involves defining system parameters such as time, security, and port settings. It also includes configuring operations such as scheduling tasks, defining approval workflows, managing licenses, managing extensible attributes, and configuring access control for supported operations.

The tasks covered in this section include:

- [Configuring Access Control](#)
- [Administrative Permissions](#)
- [Operations that Support Access Control](#)
- [Applying Access Control to Operations](#)
- [Defining Named ACLs](#)
- [Managing Named ACLs](#)
- [Managing Time Settings](#)
- [Using NTP for Time Settings](#)
- [Managing Extensible Attributes](#)
- [Configuration Examples for Inheritable Extensible Attributes](#)

- [Managing Security Operations](#)
- [Configuring Proxy Servers](#)
- [Configuring Ethernet Ports](#)
- [Using the LAN2 Port](#)
- [Using the MGMT Port](#)
- [About Lights Out Management](#)
- [Setting Static Routes](#)
- [Enabling DNS Resolution](#)
- [Managing Licenses](#)
- [About IB-FLEX](#)
- [Configuring DNS Cache Acceleration](#)
- [About Elastic Scaling](#)
- [Managing the Order of Match Lists](#)
- [Managing NIOS Appliances](#)
- [Managing the Disk Subsystem on the Infoblox-4010](#)
- [Restarting Services](#)
- [Configuring the Orphan Mode](#)

## Configuring Access Control

To effectively manage your core network services, you can grant legitimate hosts access to specific tasks and operations using an access control list (ACL) or anonymous access control entries (ACEs). Depending on your admin permissions, you can configure a named ACL, and then apply it to multiple operations, such as file distribution and DNS zone transfers. For information about admin permissions, see [About Administrative Permissions](#).

When you define a named ACL, you add access control types such as IPv4 and IPv6 addresses, IPv4 and IPv6 networks, nested named ACLs, and TSIG key based ACEs to a list, and then grant each entry in the list the Allow or Deny permission. For information about named ACLs and how to configure them, see [Defining Named ACLs](#). Note that each operation supports specific access control types. You cannot apply a named ACL to an operation that does not support the access control types contained in the named ACL. For more information about which NIOS operations support access control and which access control types each operation supports, see [Operations that Support Access Control](#).

When you add or modify a named ACL, or when you import named ACLs and ACEs to an existing named ACL through CSV import, the appliance does not automatically validate the ACEs in the list. For more information about how to import named ACLs and ACEs, refer to the [Infoblox CSV Import Reference](#). To avoid conflicts and unexpected results, you must perform validations for all named ACLs before you use them for access control. When the appliance detects a conflict or an optimized issue about a specific ACE during the validation process, it displays detailed information in a CSV file. For more information about ACL validation, see [Validating Named ACLs](#).

## Administrative Permissions

You can configure a named ACL at the Grid level and override it at the member and object level. Superusers and limited-access users with Read/Write permission to **All Named ACLs** can create, modify, and delete named ACLs. Users with Read-only permission to **All Named ACLs** can apply a named ACL to a supported object if they have Read/Write permission to the respective object. Other users can only view named ACLs and their entries. For information about admin permissions, see [About Administrative Permissions](#).

## Operations that Support Access Control

On the appliance, only certain operations support access control. You can apply one named ACL or multiple anonymous ACEs to each operation. However, you cannot apply multiple named ACLs or use a combination of named ACLs and



ACEs. Note that each operation supports different access control types. For example, DNS zone transfers support IPv4 and IPv6 addresses and networks as well as TSIG key based ACEs, while AAAA filtering supports only IPv4 addresses and networks.

When you apply a named ACL to an operation, the appliance validates to ensure that the named ACL contains ACEs that are supported by the operation. The appliance also validates any new ACEs that you add to an existing named ACL. If a named ACL contains access control types that an operation does not support, the appliance displays an error message and you cannot apply that named ACL to the operation. Thus when defining a named ACL for a specific operation or applying an existing named ACL, ensure that it contains access control types that the operation supports. The below table lists access control types for NIOS operations that support access control.

*Operations that Support Access Control*

Operation Type	Supported Access Control Types				
	IPv4 Addresses and Networks	IPv6 Addresses and Networks	TSIG Key Based ACEs	DNSone 2.x TSIG Key	Any Address and Network
GUI and API Access	Yes	Yes	No	No	No
NTP Service and NTP Queries	Yes	Yes	No	No	Yes
File Distribution Services	Yes	No	No	No	No
Syslog Proxy Access Control	Yes	Yes	No	No	No
DNS Zone Transfers (excludes zone transfers for Microsoft servers)*	Yes	Yes	Yes	Yes	Yes
Dynamic DNS Updates	Yes	Yes	Yes	No	Yes
DNS Queries	Yes	Yes	Yes	No	Yes
Recursive Queries	Yes	Yes	Yes	No	Yes
Blackhole Lists	Yes	Yes	No	No	Yes
AAAA Filtering	Yes	No	No	No	Yes
Forward DNS Updates	Yes	Yes	Yes	No	Yes
Match Clients for DNS Views	Yes	Yes	Yes	Yes	Yes
Match Destinations for DNS Views	Yes	Yes	Yes	Yes	Yes
DNS64 Clients	Yes	Yes	No	No	Yes

Operation Type	Supported Access Control Types				
	IPv4 Addresses and Networks	IPv6 Addresses and Networks	TSIG Key Based ACEs	DNSone 2.x TSIG Key	Any Address and Network
DNS64 Mapped	Yes	No	No	No	Yes
DNS64 Exclude IPv6	No	Yes	No	No	Yes



#### Note

\* Zone transfers for Microsoft servers do not support named ACLs. However, you can still use individual ACEs to configure access control. For more information about how to configure zone transfer settings for Microsoft servers, see [Setting Zone Properties](#). In addition, the DNSone 2.x TSIG key supports only the "Allow" permission. You cannot change "Allow" to "Deny."

Complete the following tasks to use a named ACL:

1. Define a named ACL, as described in [Defining Named ACLs](#).
2. Validate the named ACL, as described in [Validating Named ACLs](#).
3. Apply the named ACL to specific operations, as described in [Applying Access Control to Operations](#).

## Applying Access Control to Operations

When you apply access control to NIOS operations, you can use anonymous ACEs or a named ACL. You cannot combine ACEs and named ACLs for access control. Depending on the access control types each operation supports, you may or may not be able to apply a named ACL to a specific operation. For information about which access control types each operation supports in *Operations that Support Access Control* table, see [Operations that Support Access Control](#).



#### Note

If you disable access control or select **None** or **Any** for an operation, the appliance removes the previously applied named ACL or the configured anonymous ACEs. To avoid losing your ACE configuration, Infoblox recommends that you convert the ACEs to a named ACL.

For information about how to apply access control to each supported operation, see the following:

- DNS zone transfers, as described in [Enabling Zone Transfers](#).
- DNS queries, as described in [Controlling DNS Queries](#).
- Recursive queries, as described in [Enabling Recursive Queries](#).
- Dynamic DNS updates, as described in [Configuring DNS Servers for DDNS](#).
- AAAA filtering, as described in [Controlling AAAA Records for IPv4 Clients](#).
- Blackhole list, as described in [Configuring a DNS Blackhole List](#).
- Match clients list for DNS views, as described in [Defining Match Clients Lists](#).
- Match destinations for DNS views, as described in [Defining a Match Destinations List](#).
- DNS64 clients, DNS64 mapped IPv4 addresses, and DNS64 excluded IPv6 addresses, as described in [Setting DNS64 Group Properties](#).
- File distribution services, as described in [Configuring Access Control for File Distribution](#).

- Grid Manager and API access, as described in [Configuring Security Features](#).
- NTP access control, as described in [Defining NTP Access Control](#).
- Syslog proxy access, as described in [Configuring Syslog for Grid Members](#).

## Defining Named ACLs

Depending on how you plan to use a named ACL and which access control types an operation supports, you can add one or all of the following when you define a named ACL: IPv4 and IPv6 addresses, IPv4 and IPv6 networks, TSIG key based ACEs, DNSone 2.x TSIG keys. You can also add an existing named ACL as a nested ACL to a new or existing named ACL.

When configuring a named ACL, ensure that you define it correctly for the intended operations using the supported access control types. For example, if you want to apply a named ACL to AAAA filtering, do not include IPv6 addresses or networks in the named ACL because AAAA filtering does not support IPv6 addresses and networks. For information about which access control types in *Operations that Support Access Control* table, see [Operations that Support Access Control](#).

To define a named ACL:

1. From the **Administration** tab, select the **Named ACLs** tab, and then click the Add icon.
2. In the *Add Named ACL* wizard, complete the following:
  - **Name:** Enter a name for the named ACL. You can enter up to 64 characters.
  - **Comment:** Enter additional information about the named ACL.
3. Click **Next**. Complete the following to add ACEs to the named ACL:
  - Click the Add icon and select one of the following access control types from the drop-down list. Depending on your selection, Grid Manager adds a row to the table directly or expands the panel before adding a row.
    - **IPv4 Address:** Select this to add an IPv4 address. Click the **Entry** field and enter the IPv4 address. The **Operation** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
    - **IPv4 Network:** When you select this, enter the network address in the **Address** field, select the netmask using the slider, and then select **Allow** or **Deny** from the **Permission** drop-down list. Click **Add** and Grid Manager adds the entry to the table.
    - **IPv6 Address:** Select this to add an IPv6 address. Click the **Entry** field and enter the IPv6 address. The **Operation** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
    - **IPv6 Network:** When you select this, enter the network address and its netmask in the **Address** field, and then select **Allow** or **Deny** from the **Permission** drop-down list. Click **Add** and Grid Manager adds the entry to the table.
    - **TSIGKey:** In the **AddTSIGKey** panel, complete the following, and then click **Add** to add the TSIG key to the list:
      - **Key name:** Enter a meaningful name for the key, such as a zone name or the name of the remote name server. This name must match the name of the same TSIG key on other name servers.  
Note that the **Key name** must be a valid domain name without any space and it cannot begin with DHCP\_UPDATER.
      - **Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.
      - **Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.
    - **DNSone 2.x TSIG Key:** Select this when the client is a NIOS appliance running DNS One 2.x code. The appliance automatically populates the value of the key in the **Entry** field. The **Operation** column displays **Allow** by default. You cannot change the default permission.
    - **Any Address/Network:** Select this to allow or deny permission for any addresses and networks.
    - **Named ACL:** When you select this, Grid Manager displays the *Named ACLs* Selector. Select the named ACLs you want to add to the new ACL. If you have only one existing named ACL, Grid Manager automatically adds the named ACL to the list. The selected named ACL becomes a nested ACL in the newly created named ACL.

Note: The **Order** field in the table displays the position of each entry based on the order it is placed in the list. You can modify this number to change the order of an ACE. You can also select the ACE checkbox and use the up and down arrows next to the table to place the entry in the desired position.

4. Click **Next** to enter extensible attributes for the named ACL. For information, see [About Extensible Attributes](#).
5. Save the configuration.

## Managing Named ACLs

You can do the following after you have configured named ACLs for access control:

- Preview the list of ACEs in a named ACL, as described in [Previewing ACEs in Named ACLs](#) below.
- Validate ACEs in a named ACL, as described in [Validating Named ACLs](#) below.
- View a complete list of configured named ACLs, as described in [Viewing Named ACLs](#) below.
- Modify information in a named ACL, as described in [Modifying Named ACLs](#) below.
- Apply a named ACL to supported operations, as described in [Applying Access Control to Operations](#).
- Delete a named ACL, as described in [Deleting Named ACLs](#) below.
- Export and print the list of named ACLs.

### Previewing ACEs in Named ACLs

You can preview the list of ACEs in a named ACL when you add or modify it. When you click the Preview icon in the *Add Named ACL* wizard or *Named ACL* editor, the appliance lists all the entries in the named ACL, even if you have selected only one or a few entries in the wizard or editor.

To preview a named ACL:

1. From the **Administration** tab, select the **Named ACLs** tab -> *named\_acl* checkbox, and then click the Preview icon.
2. In a separate browser window, Grid Manager displays the following information for each ACE in the named ACL:
  - **Entry:** Displays one of the following: IPv4 or IPv6 address, IPv4 or IPv6 network, or TSIG key. Note that if the named ACL contains nested ACLs, all entries in the nested ACLs are displayed in a flat view. Grid Manager does not display the name of the nested ACL.
  - **Type:** The access control type of the entry. This can be **IPv4 Address**, **IPv6 Address**, **IPv4 Network**, **IPv6 Network**, **TSIG Key**, or **DNSone 2.x TSIG Key**.
  - **Operation:** Displays the access permission for the entry. This can be **Allow** or **Deny**.

### Validating Named ACLs

When you add or modify a named ACL, the appliance does not automatically validate the ACEs in the list. In addition, when you import named ACLs or ACEs to a named ACL, no automatic validation is performed. To avoid unintended consequences, ensure that you validate your named ACLs before you save them or use them for access control.



#### Note

When you click the Validate icon in the *Add Named ACL* wizard or *Named ACL* editor, the appliance validates all the entries in the named ACL, even if you have selected only one or a few entries in the wizard or editor.

The following examples demonstrate the importance of validating named ACLs:

#### Example 1

You configure a named ACL "foo" that includes a Deny permission to 10.0.0.0/16. You then assign "foo" to DNS zone transfers. You later import an "Allow/10.0.0.0/24" entry to "foo" through CSV import. The appliance appends the entry to the end of "foo." When you perform an ACL validation on "foo" after a DNS service restart, the appliance displays a warning message indicating that the new "Allow/10.0.0.0/24" entry is now included in the previously configured "Deny/10.0.0.0/16" entry. Since DNS service works on a first-match access control basis, zone transfers will not be allowed for the 10.0.0.0/24 network, which is probably not your original intent. You can then modify the named ACL to correct this error. On the other hand, if you do not perform the ACL validation, the appliance is not notified about the new network entry in "foo." As a result, you are not notified about the denial of zone transfers to 10.0.0.0/24.

## Example 2

You add a nested named ACL "bar" as the first entry to the named ACL "foo." You then add a "Deny All" entry right after "bar" (as the second entry). You later import a new "Allow All" entry to "bar" through CSV import. The "Allow All" entry will be appended to the end of "bar." When you perform an ACL validation on "foo" after the CSV import, the appliance detects a conflict between the "Allow All" (in "bar") and "Deny All" (right after "bar") permissions and displays a warning. Imagine if you do not perform the ACL validation on "foo," the appliance is not notified about the new "Allow All" entry in "bar" and therefore cannot detect the conflict between the "Allow All" and "Deny All" entries. As a result, almost all hosts will get zone transfers, which may not be the outcome you have intended.



### Note

It is important that you manually validate each named ACL after a CSV import to ensure data and performance integrity. The appliance does not automatically perform ACL validation.

To validate a named ACL:

1. From the **Administration** tab, select the **Named ACLs** tab -> `named_acl` checkbox, and then click the Validate icon.  
or  
In the *Add Named ACL* wizard or *Named ACL* editor, click the Validate icon.
2. Grid Manager validates all the ACEs in the named ACL and displays a system message at the top of the screen indicating whether all ACEs in the named ACL are valid or not, depending on the validation results. When the appliance detects conflicts or issues related to specific ACEs, it displays the results in a CSV file. You can save the file or open it. Grid Manager displays the following information in the file:
  - **Defined ACL:** The name of the named ACL.
  - **Type of Issue:** The type of issue found. This can be one of the following:
    - **Optimize:** An ACE is a duplicate of a previous entry or an ACE configuration can be a subset of another entry. See optimized suggestions in the **Issue** field.
    - **Conflict:** The same IP address or network has a conflicting permission. Re-configure the ACE based on your requirements.
    - **Warning:** An ACE is a subset of a previously configured entry, but it has a conflicting permission.
  - **ACE A:** The ACE that has a conflict or an optimized issue with ACE B.
  - **ACE B:** The ACE that has a conflict or an optimized issue with ACE A.
  - **Issue:** Detailed information and optimized suggestions about the conflict or issue.



### Note

It may take a long time to validate a named ACL that contains a large number of ACEs.

## Viewing Named ACLs

To view a list of named ACLs:

- From the **Administration** tab, select the **Named ACLs** tab. Grid Manager displays the following information for each named ACL:
  - **Name:** The name of the named ACL.
  - **Comment:** Information about the named ACL.
  - **Site:** The site to which the named ACL belongs. This is one of the predefined extensible attributes.

You can also do the following in the **Named ACLs** tab:

- Modify data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes or **Cancel** to exit.
- Sort the named ACLs in ascending or descending order by column.
- Select a named ACL and click the Edit icon to modify data, or click the Delete icon to delete it.

- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the Go to field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Print and export data in this tab.

## Modifying Named ACLs

You can modify ACEs in an existing named ACL. When you update a named ACL, the appliance validates the updates to ensure that ACEs in the named ACL are valid for the operations to which the named ACL has been applied. For example, if a named ACL is used for file distribution access, you are not allowed to add IPv6 address access control to it because the file transfer operation does not support IPv6 addresses.

To modify a named ACL:

1. From the **Administration** tab, select the **Named ACLs** tab -> *named\_acl* checkbox, and then click the Edit icon.
2. The *Named ACL* editor provides the following tabs from which you can modify data:
  - **General Basic:** You can modify data in this tab as described in [Defining Named ACLs](#).
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific named ACL. You can also modify the values of the extensible attributes. For information, see [About Extensible Attributes](#).
  - **Permissions:** This tab appears only if you belong to a superuser admin group. For information about managing permissions, see [About Administrative Permissions](#).

## Deleting Named ACLs

When you delete a named ACL, the appliance puts it in the Recycle Bin, if enabled. You can restore the named ACL later if needed.



### Note

You cannot delete a named ACL that has been applied to an operation or is currently in use by another operation.

To delete a named ACL:

1. From the **Administration** tab, select the **Named ACLs** tab -> *named\_acl* checkbox, and then click the Delete icon. You can select multiple named ACLs for deletion.
2. In the *Delete Confirmation* dialog box, click **Yes**.

## Managing Time Settings

You can define the date and time settings for your NIOS appliance using the Infoblox Appliance Startup Wizard. Alternatively, you can set the date and time of the appliance anytime after you first configure it if you did not do so using the startup wizard or if you need to change it if, for example, you move an appliance from a location in one time zone to a location in a different time zone. To set the date and time of the appliance, you can either manually enter the values or configure the appliance to synchronize its time with a public NTP server.

### Changing Time and Date Settings

If you do not use the NTP service, you can set the date and time for a Grid.



### Note

You cannot manually set the date and time if the NTP service is enabled.

To set the time and date for a Grid using the *Grid Properties* editor:

1. From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.
2. In the **General** tab of the *Grid Properties* editor, complete the following:
  - **Date:** Click the calendar icon to select a date or enter the date in YYYY-MM-DD format.
  - **Time:** Click the clock icon to select a time or enter the time in HH:MM:SS format. For afternoon and evening hours, use the integers 13-24.
3. Save the configuration and click **Restart** if it appears at the top of the screen.



#### Note

Changing the date and time resets the application and terminates the management session.

## Changing Time Zone Settings

Whether you enable NTP (Network Time Protocol) or manually configure the date and time, you must always set the time zone manually. You can set the time zone for a Grid, which then applies to all members. If different members are in different time zones, you can choose the time zone that applies to most members at the Grid level, and then override the setting for the remaining members.



#### Note

Changing the time zone does not reset the application nor does it terminate the management session.

To set the time zone for a Grid or member:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.  
**Member:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the **General** tab of the editor, select the appropriate time zone.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Monitoring Time Services

In a Grid, the Grid Master and its members use an internal NTP daemon to synchronize their time. It is not user-configurable and functions regardless of how you set the time on the Grid Master. The *Detailed Status* panel contains an NTP Synchronization icon so you can monitor the internal NTP daemon that runs within a Grid to ensure the time among its members is synchronized.

To display the *Detailed Status* panel, from the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Detailed Status icon in the table toolbar of the *Members* panel. The following are descriptions of the NTP status icons in the *Detailed Status* panel:

Icon	Color	Meaning
	Green	The NTP service is running properly.
	Yellow	The appliance is synchronizing its time.
	Red	The NTP service is not running properly. View the corresponding description for additional information.

## Using NTP for Time Settings

NTP (Network Time Protocol) is a standard protocol that system clocks use to ensure their time is always accurate. Appliances that use NTP try to get their time as close as possible to UTC (Coordinated Universal Time), the standard timescale used worldwide. NTP uses UDP (User Datagram Protocol) on port 123 for communications between clients and servers.

NTP is based on a hierarchy where reference clocks are at the top. Reference clocks use different methods such as special receivers or satellite systems to synchronize their time to UTC. NTP servers on the first level of the hierarchy synchronize their time with the reference clocks, and serve time to clients as well. Each level in the hierarchy is a stratum; stratum-0 is a reference clock. Stratum-1 servers synchronize their clocks with reference clocks. Stratum-2 servers synchronize their clocks with stratum-1 servers, and so forth. The stratum number indicates the number of levels between the NTP server and the reference clock. A higher stratum number could indicate more variance between the NTP server and the reference clock.

You can configure a NIOS appliance to function as an NTP client that synchronizes its clock with an NTP server. For more information, see [NIOS Appliances as NTP Clients](#) below. NTP clients typically use time information from at least three different sources to ensure reliability and a high degree of accuracy. There are a number of public NTP servers on the Internet with which the NIOS appliance can synchronize its clock. For a list of these servers, you can access <http://www.ntp.org>. When NTP is configured, it listens on all interfaces, including the loopback interface on the NIOS appliance.

In a Grid, the Grid Master and Grid members can function as NTP clients that synchronize their clocks with external NTP servers. They can in turn function as NTP servers to other appliances in the network. For more information, see [NIOS Appliances as NTP Servers](#). Note that when the Grid Master functions as an NTP server, it synchronizes its local clock with its NTP clients and does not synchronize time with any other external NTP server. This allows you to deploy multiple NTP servers to ensure accurate and reliable time across the network. To configure the Grid Master and Grid members as NTP clients, you must first enable the NTP service and configure external NTP servers at the Grid level. You can then configure the Grid Master and Grid members to override the Grid-level NTP servers and use their own external NTP servers. Note that a Grid member will not function as an NTP client if you do not enable the NTP service at the Grid level. A Grid member synchronizes its clock with the Grid Master if you do not configure it to use external NTP servers. When an NTP service-enabled member goes offline from the grid, that member sets the Grid Master's NTP configurations. If the Grid Master goes offline (because of a shutdown or a disconnecting network, and so on), the Grid Master Candidate and Grid members synchronize with the external NTP servers.

In case of leap second insertion, the Infoblox Grid handles the leap second over a period of time instead of performing a one-time adjustment. In other words, when using the Grid as the NTP server, it follows the standard NTP recovery process by slewing over a certain period of time when handling the leap second. The slewing process could therefore cause synchronization issues among NTP clients. The out-of-sync state is usually resolved when all NTP clients catch up with the server.

The following figure illustrates how NIOS appliances (the Grid Master and Grid members) in a Grid function as the NTP server or the NTP client, depending on your NTP configuration.

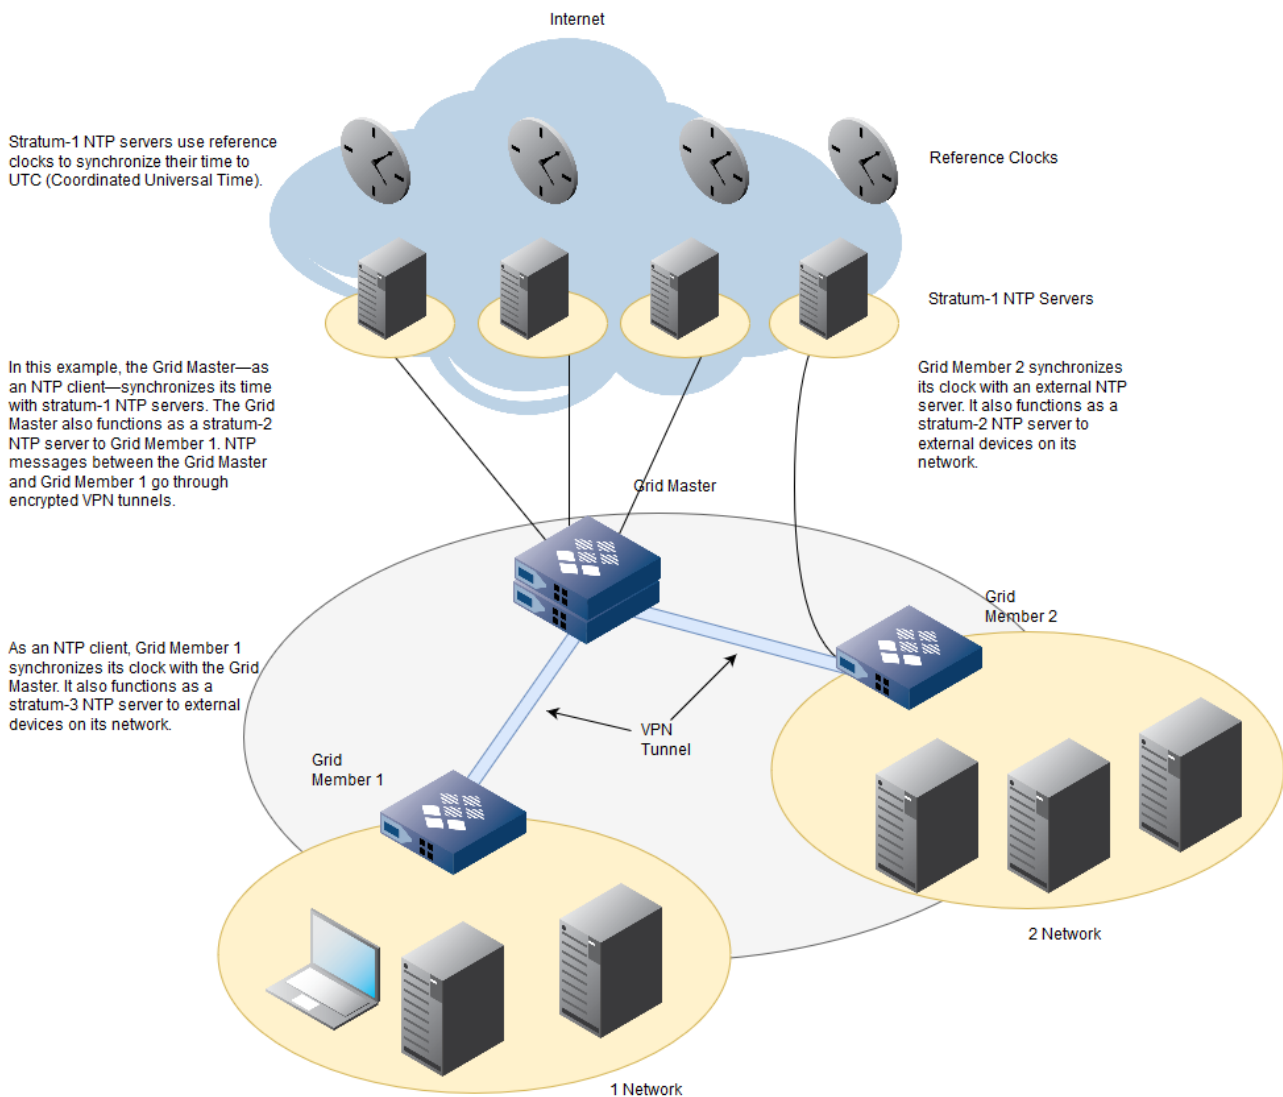


### Note

The NTP service supports both IPv4 and IPv6 networks.

*Infoblox Appliances as NTP Servers*





## Authenticating NTP

To prevent intruders from interfering with the time services on your network, you can authenticate communications between a NIOS appliance and a public NTP server, and between a NIOS appliance and external NTP clients. NTP communications within the Grid go through an encrypted VPN tunnel, so you do not have to enable authentication between members in a Grid.

NTP uses symmetric key cryptography, where the server and the client use the same algorithm and key to calculate and verify a MAC (message authentication code). The MAC is a digital thumbprint of the message that the receiver uses to verify the authenticity of a message.

As shown in the figure below, the NTP client administrator must first obtain the secret key information from the administrator of the NTP server. The server and the client must have the same key ID and data. Therefore, when you configure the NIOS appliance as an NTP client and want to use authentication, you must obtain the key information from the administrator of the external NTP server and enter the information on the NIOS appliance. When you configure a NIOS appliance as an NTP server, you must create a key and send the key information to clients in a secure manner. A key consists of the following:

- Key Number: A positive integer that identifies the key.
- Key Type: Specifies the key format and the algorithm used to calculate the MAC (message authentication code) of a message.
  - M: The key is a 1-31 character ASCII string using MD5 (Message Digest).

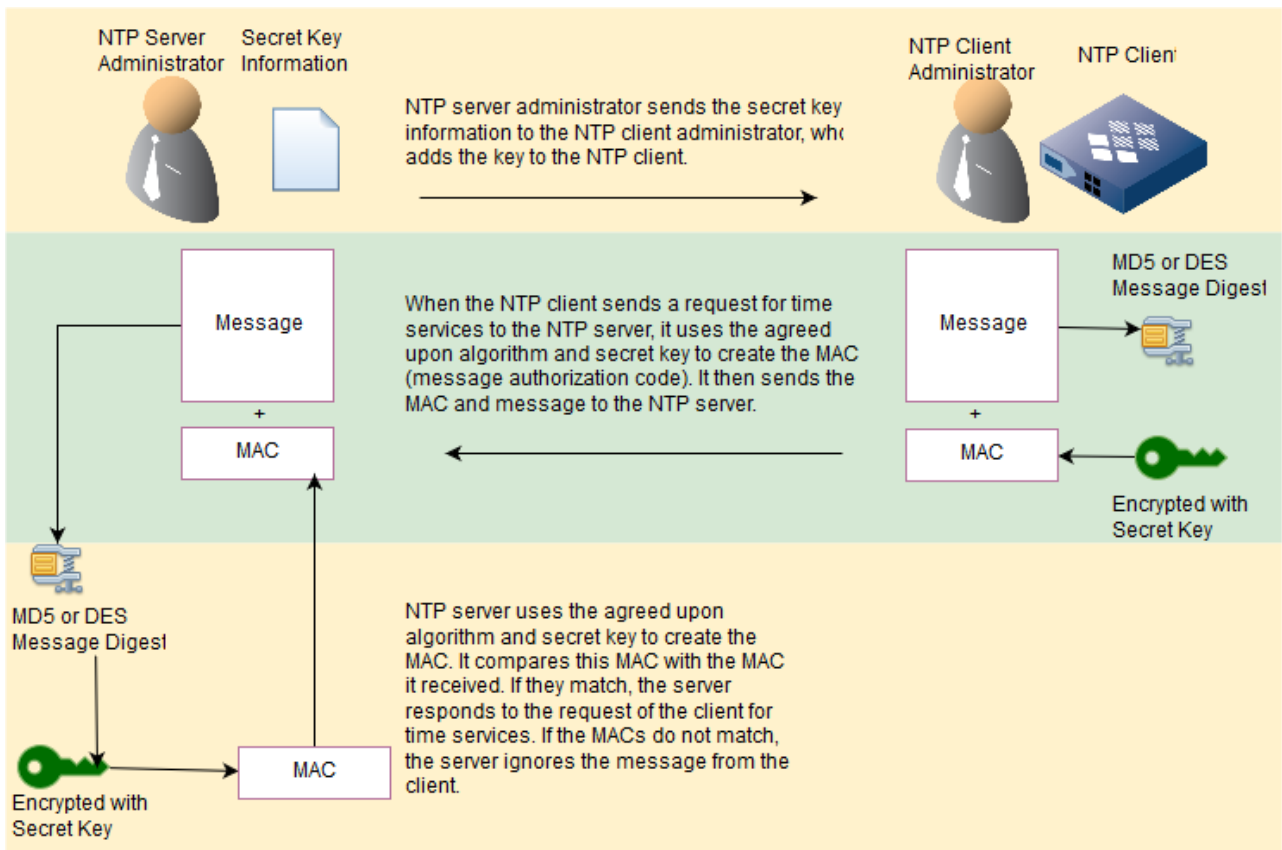
- S: The key is a 64-bit hexadecimal number in DES (Data Encryption Standard) format. The high order 7 bits of each octet form the 56-bit key, and the low order bit of each octet is given a value so that the octet maintains odd parity. You must specify leading zeros so the key is exactly 16 hexadecimal digits long and maintains odd parity.
- A: The key is a DES key written as a 1-8 character ASCII string.
- N: The key is a 64-bit hexadecimal number in NTP format. It is the same as the S format, but the bits in each octet have been rotated one bit right so the parity bit is in the high order bit of the octet. You must specify leading zeros and odd parity must be maintained.
- Key String: The key data used to calculate the MAC. The format depends on the Key Type you select.

When the NTP client initiates a request for time services to the NTP server, it creates the MAC by using the agreed upon algorithm to compress the message and then encrypts the compressed message (which is also called a message digest) with the secret key. The client appends the MAC to the message it sends to the NTP server. When the NTP server receives the message from the client, it performs the same procedure on the message — it compresses the message it received, encrypts it with the secret key and generates the MAC. It then compares the MAC it created with the MAC it received. If they match, the server continues to process and respond to the message. If the MACs do not match, the receiver drops the message.

The following table lists the NTP client server behaviour in different scenarios:

Scenario	Behavior
No authentication on both the NTP client and server	The NTP client will synchronize with the server
Authentication on the NTP server, no authentication on the NTP client	The NTP client will synchronize with the server
Authentication on both the NTP server and client	The NTP client will synchronize with the server
No authentication on the NTP server, authentication on the client	The NTP client will be out-of-synchronization with the server

*NTP Client Administrator Obtaining Secret Key from NTP Server Administrator*



### NIOS Appliances as NTP Clients

You can configure an independent NIOS appliance, a Grid Master, or any Grid member in a Grid as an NTP client that synchronizes its system clock with an external NTP server.



#### Note

You can configure NIOS appliance as NTP client in either IPv4, IPv6, or dual mode (IPv4 and IPv6) network environment.

When you enable a NIOS appliance to function as an NTP client, you must specify at least one NTP server with which the appliance can synchronize its clock. Infoblox recommends that you specify multiple NTP servers that synchronize their time with different reference clocks and that have different network paths. This increases stability and reduces risk in case a server fails. For a list of public NTP servers, you can access [www.ntp.org](http://www.ntp.org).

When you specify multiple NTP servers, the NTP daemon on the appliance determines the best source of time by calculating round-trip time, network delay, and other factors that affect the accuracy of the time. NTP periodically polls the servers and adjusts the time on the appliance until it matches the best source of time. If the difference between the appliance and the server is less than five minutes, the appliance adjusts the time gradually until the clock time matches the NTP server. If the difference in time is more than five minutes, the appliance immediately synchronizes its time to match that of the NTP server.

To secure communications between a NIOS appliance and an NTP server, you can authenticate communications between the appliance and the NTP server. When you configure authentication, you must obtain the key information from the administrator of the NTP server and enter the key on the appliance. For information, see [Authenticating NTP](#) above. In a Grid, you can configure the Grid Master and Grid members to synchronize their clocks with external NTP servers. When you enable the NTP service on the Grid, the Grid Master automatically functions as an NTP server to the Grid members. A Grid member can synchronize its time with the Grid Master, an external NTP server, or another Grid member. When Grid members synchronize their times with the Grid Master, the Grid Master and its members send NTP

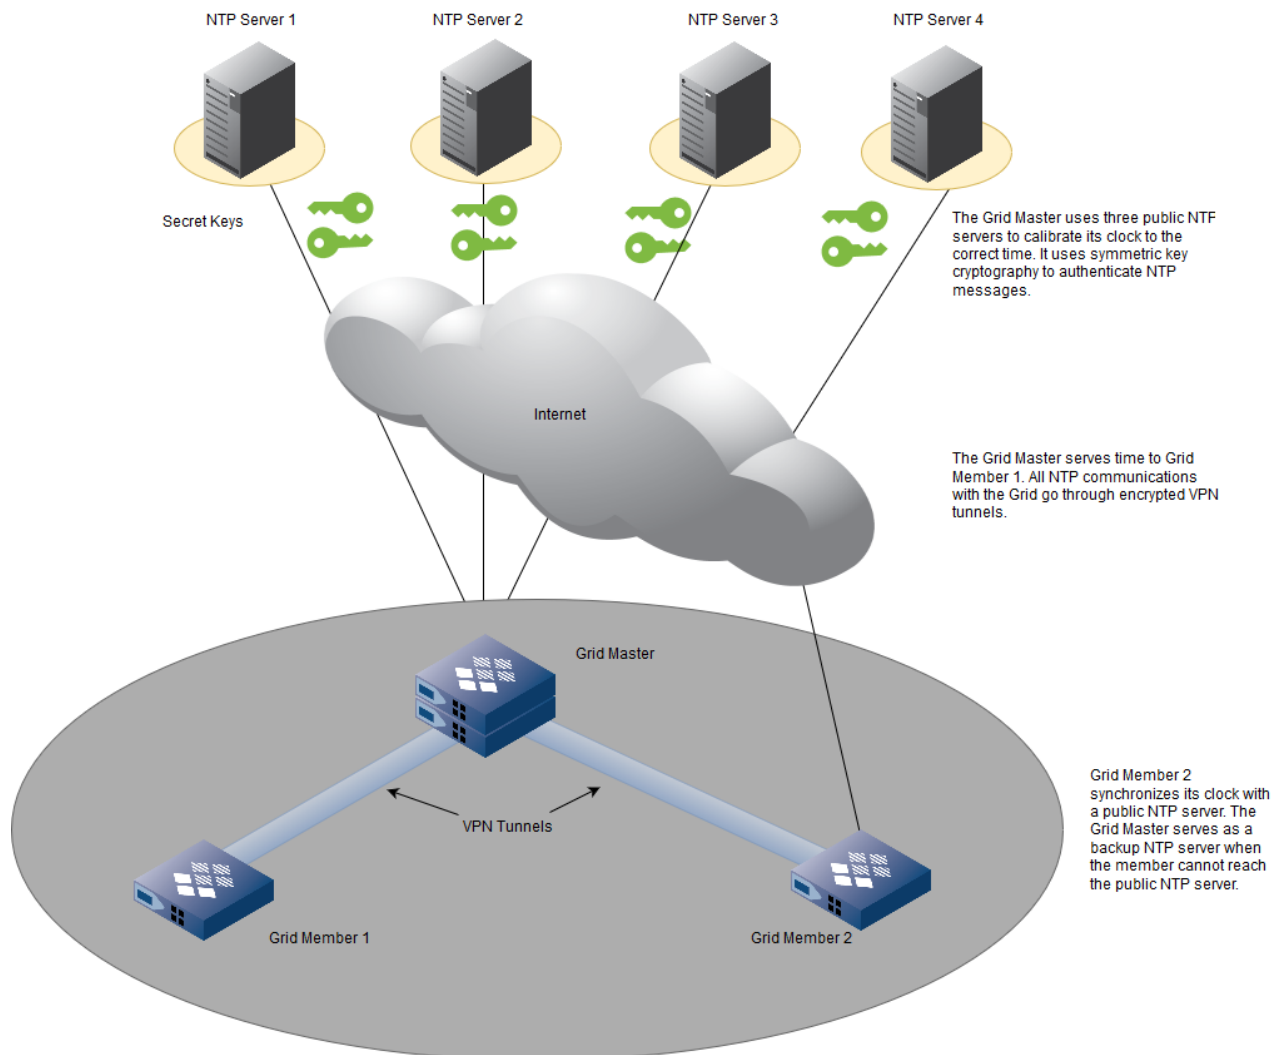
messages through an encrypted VPN tunnel, as shown in the following figure. When a Grid member synchronizes its time with another Grid member, the NTP messages are not sent through a VPN tunnel.



**Note**

Grid member cannot act as an NTP server to the Grid Master.

**Grid Master as NTP Client**



**Configuring the Grid to Use NTP**

In a Grid, the Grid Master and Grid members can synchronize their clocks with external NTP servers. They then forward the clock time to other appliances in the network. Likewise, in an independent HA pair, the active node communicates directly with an external NTP server. The passive node then synchronizes its clock with the active node.

In a Grid, you must first enable the NTP service and configure external NTP servers at the Grid level before you configure the Grid Master and Grid members as NTP clients.

To configure a Grid Master as an NTP client, perform the following tasks:

- If you want to enable authentication between the Grid members and NTP servers, you must specify the authentication keys before enabling the NTP service. You can specify authentication keys at the Grid and member levels.
- Enable the NTP service on the Grid and specify one or more external NTP servers. For information, see Synchronizing the Grid with External NTP Servers below.

### Adding NTP Authentication Keys

To enable authentication between the appliance and the NTP servers, add the authentication keys before enabling the NTP service on the Grid. You can specify authentication keys at the Grid and member levels.

To add NTP authentication keys:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **NTP -> NTP Grid Config**.  
**Member:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox. Expand the Toolbar and click **NTP -> NTP Member Config**.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. Click the Add icon in the NTP Keys section and enter the following information.
  - **Key Number:** A positive integer that identifies a key.
  - **Type:** Specifies the key format and the algorithm used to calculate the MAC (message authentication code) of a message.
    - **MD5 in ASCII format (M):** The key is a 1-31 character ASCII string using MD5 (Message Digest).
    - **DES in hex format (S):** The key is a 64-bit hexadecimal number in DES (Data Encryption Standard) format. The high order 7 bits of each octet form the 56-bit key, and the low order bit of each octet is given a value so that the octet maintains odd parity. You must specify leading zeros so the key is exactly 16 hexadecimal digits long and maintains odd parity.
    - **DES in ASCII format (A):** The key is a DES key written as a 1-8 character ASCII string.
    - **DES in NTP format (N):** The key is a 64-bit hexadecimal number in NTP format. It is the same as the S format, but the bits in each octet have been rotated one bit right so the parity bit is in the high order bit of the octet. You must specify leading zeros and odd parity must be maintained.
    - **SHA1 in ASCII format (SHA1):** The key is a 40-character hexadecimal string, and it uses hash based symmetric encryption algorithm.  
Note that FIPS compliance is not supported in SHA-1.
  - **String:** The key data used to calculate the MAC. The format depends on the Key Type you select.
3. Click **Save** to save the entry and keep the editor open so you can enable the Grid to synchronize its time with external NTP servers.

Note that if you enter a new key, the appliance checks if the key already exists in the key list. If the key exists, but either the key type or key string does not match, the NIOS appliance sends an error message.

After you enter an authentication key, you can modify or delete it. Note that you cannot delete a key that an NTP server references. You must first delete all NTP servers that reference that key and then delete the key.

### Synchronizing the Grid with External NTP Servers

To enable the Grid to synchronize its time with external NTP servers:

1. From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **NTP -> NTP Grid Config**.
2. In the **General** tab of the *Grid NTP Properties* editor, select **Synchronize the Grid with these External NTP Servers**.
3. Click the Add icon to add external NTP servers and enter the following information in the *Add NTP Server* dialog box:
  - **NTP Server (FQDN or IP Address):** Enter either the IP address or the resolvable host name of an NTP server. Entries may be an IPv4 or IPv6 address. You can view a list of public NTP servers at [ntp.isc.org](http://ntp.isc.org). To check whether the DNS server can resolve the NTP server host name, click **Resolve Name**. You must have a DNS name resolver configured. For information, see [Enabling DNS Resolution](#). As a best practice when configuring NTP servers using the FQDN (instead of an IP address), configure an external DNS name resolver that is reachable by NIOS appliance. Otherwise, the NIOS system startup can be delayed.
  - **Enable Authentication:** Select this option to enable authentication of NTP communications between the external NTP server and the NIOS appliance (the Grid Master or Grid member in a Grid, an independent NIOS appliance, or the active node in an independent HA pair).

Note that to prevent intruders from interfering with the time services on your network, you can authenticate communications between a Grid member and an external NTP server, as well as between a Grid member and external NTP clients. NTP communications within the Grid go through an encrypted VPN tunnel, so you do not have to enable authentication between the Grid Master and Grid members.

**Authentication Key:** Select a key that you previously entered from the drop-down list.

- Click **Add** to add the NTP server to the list or **Cancel** to cancel the operation. In the table, you can configure some of the following settings:
    - **Preferred:** Select this to mark an external NTP server as the preferred NTP server. You can select only one server as the preferred NTP server. NIOS uses the responses from this preferred server over responses from other external NTP servers. A response from a preferred server will be discarded if it differs significantly from the responses of other servers. Infoblox recommends that you select an NTP server that is known to be highly accurate as the preferred server, such as one that has special time monitoring hardware. Note that this option is enabled only when you have selected the checkbox **Synchronize the Grid with these External NTP Servers**.
    - **Server:** Displays the FQDN or IP address of the NTP server that you added.
    - **Authentication:** When you enable authentication, this column displays **Yes**. Otherwise, it displays **No**.
    - **Key Number:** Displays the authentication key that you have selected.
    - **BURST:** Select this checkbox to configure the NTP client to send a burst of eight packets if the external NTP server is reachable and a valid source of synchronization is available. The NTP client transmits each packet at a regular interval of two seconds. After you add an NTP server and save the configuration, the appliance will enable this option by default. When you deselect this checkbox, the client sends a single packet only once to the server.
    - **IBURST:** Select this checkbox to configure the NTP client to send a burst of eight packets if the external NTP server is not reachable when the client sends the first packet to the server. The NTP client transmits each packet at a regular interval of two seconds. After you add an NTP server and save the configuration, the appliance will enable this option by default. When you deselect this checkbox, the client sends a single packet only once to the server.
- For information about adding NTP authentication keys, see Adding NTP Authentication Keys above.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring Grid Members to Use NTP

Once you configure a Grid member to use external NTP server, make sure that the NTP service is enabled at Grid level. Otherwise, the Grid member will not function as an NTP client.

To configure Grid members to synchronize their time with external NTP servers:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox.
2. Expand the Toolbar and click **NTP** -> **NTP Member Config**.
3. In the **General** tab of the *Member NTP Properties* editor, do the following:
  - **Enable the NTP Server on this Member:** Select this checkbox to configure a Grid Master or a Grid member as an NTP server. If you have configured DNS anycast on the appliance, it can answer NTP requests through the anycast IP address.
  - **Synchronize this Member only with the Grid Master:** Select this checkbox to enable this Grid member to synchronize its time with the Grid Master. This is the default.
  - **Synchronize this Member with other NTP Servers:** Select this checkbox to enable this Grid member to use external NTP servers. When you select this checkbox, you must enter at least one external NTP server for the member.
  - **Exclude the Grid Master as an NTP Server:** Select this checkbox if you want to exclude the Grid Master from being one of the time sources. By default, the appliance automatically configures the Grid Master as the backup NTP server for a Grid member. When the member cannot reach any of its configured NTP servers, it uses the Grid Master as the NTP server. The appliance does not display the Grid Master in the NTP external server list. For a Grid Master, this checkbox has no meaning.
  - **External NTP Servers:** Click **Override** and then click the **Add** icon. In the *Add NTP Server* dialog box, enter the following information:
    - **NTP Server (FQDN or IP Address):** Enter either the IP address or the resolvable host name of an NTP server. You can view a list of public NTP servers at [ntp.isc.org](http://ntp.isc.org). To check whether the DNS server can resolve the NTP server host name, click **Resolve Name**. You must have a DNS name resolver configured.

- **Enable Authentication:** Select this checkbox to enable authentication of NTP communications between the external NTP server and the NIOS appliance (the Grid Master or Grid member in a Grid, an independent NIOS appliance, or the active node in an independent HA pair).  
Note that to prevent intruders from interfering with the time services on your network, you can authenticate communications between a Grid member and an external NTP server, as well as between a Grid member and external NTP clients. NTP communications within the Grid go through an encrypted VPN tunnel, so you do not have to enable authentication between the Grid Master and Grid members.  
**Authentication Key:** Select a key that you previously entered from the drop-down list. Note that you must enter authentication keys at the Grid level when you configure a Grid Master or Grid member to use external NTP servers.
- Click **Add** to add the NTP server to the list or **Cancel** to cancel the operation. In the table, click **Override** in the table to override configurable settings. To inherit the same properties as the Grid, click **Inherit**.
  - **Preferred:** Select this to mark an external NTP server as the preferred NTP server. You can select only one server as the preferred NTP server. NIOS uses the responses from this preferred server over responses from other external NTP servers. A response from a preferred server will be discarded if it differs significantly from the responses of other servers. Infoblox recommends that you select an NTP server that is known to be highly accurate as the preferred server, such as one that has special time monitoring hardware. Note that this option is enabled only when you have selected the checkbox **Synchronize this Member with other NTP Servers**.
  - **Server:** Displays the FQDN or IP address of the NTP server that you added.
  - **Authentication:** When you enable authentication, this column displays **Yes**. Otherwise, it displays **No**.
  - **Key Number:** Displays the authentication key that you have selected.
  - **BURST:** Select this checkbox to configure the NTP client to send a burst of eight packets if the external NTP server is reachable and a valid source of synchronization is available. The NTP client transmits each packet at a regular interval of two seconds. After you add an NTP server and save the configuration, the appliance will enable this option by default. When you deselect this checkbox, the client sends a single packet only once to the server.
  - **IBURST:** Select this checkbox to configure the NTP client to send a burst of eight packets if the external NTP server is not reachable when the client sends the first packet to the server. The NTP client transmits each packet at a regular interval of two seconds. After you add an NTP server and save the configuration, the appliance will enable this option by default. When you deselect this checkbox, the client sends a single packet only once to the server.  
Note that NTP members inherit NTP properties from the Grid. Click **Override** in the *Member NTP Properties* wizard to override configurable settings. To inherit the same properties as the Grid, click **Inherit**. For information about adding NTP authentication keys, see Adding NTP Authentication Keys above.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Managing External NTP Servers

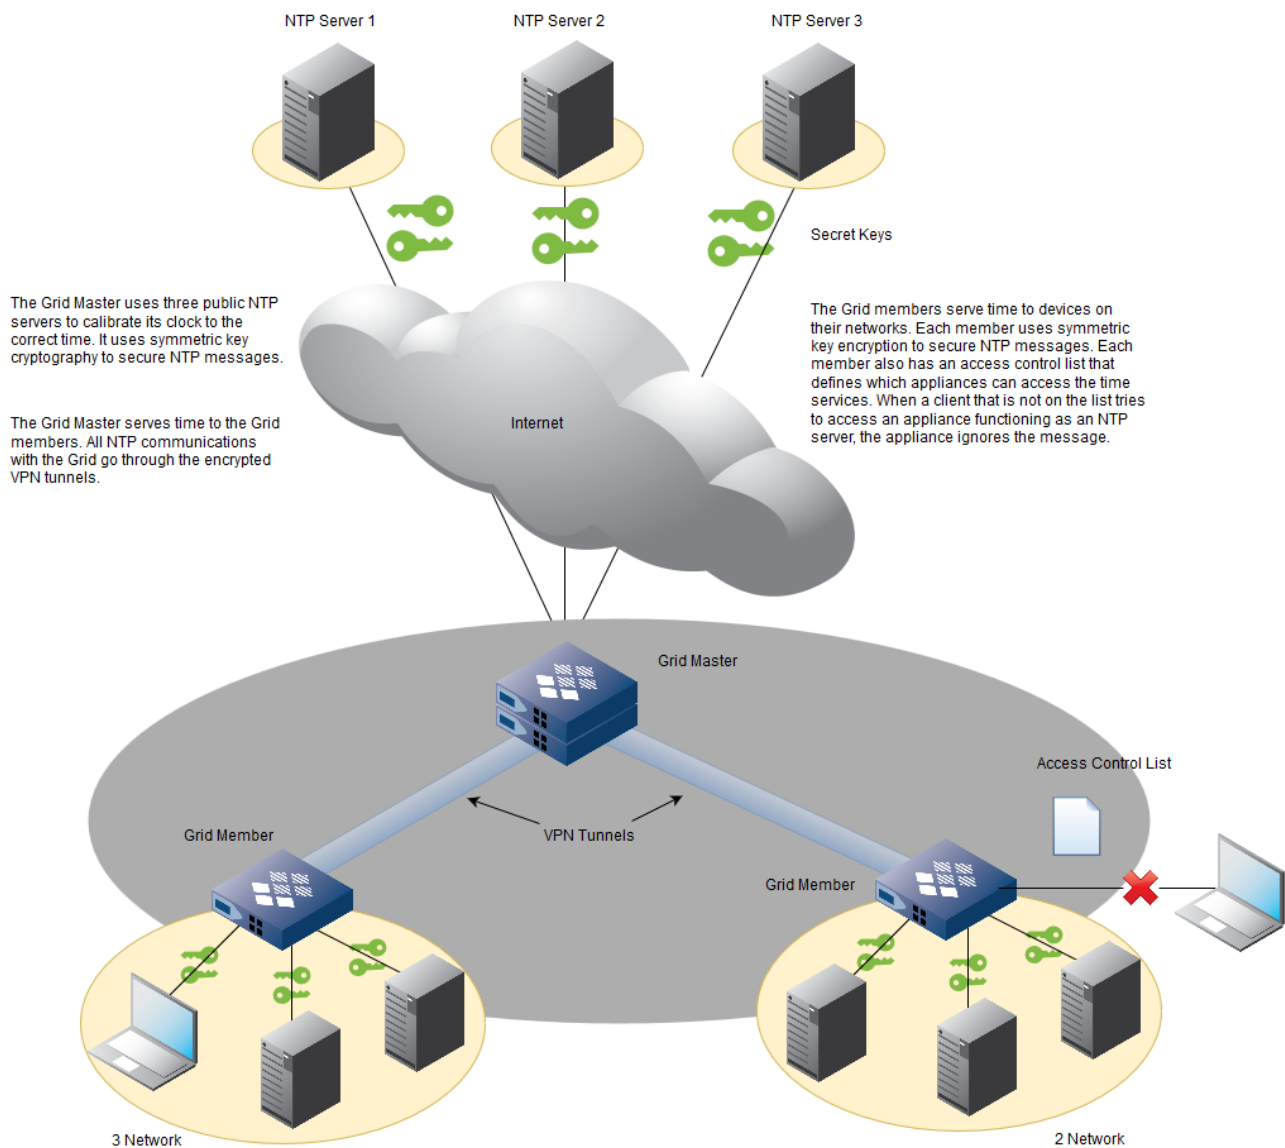
You can specify multiple NTP servers for failover purposes. The NIOS appliance attempts to connect to the NTP servers in the order they are listed. A Grid member uses the Grid Master as the NTP server when it cannot reach any of its external NTP servers.

You can change the order of the list by selecting an NTP server and dragging it to its new location or by clicking the up and down arrows. You can add and delete servers and modify their information as well.

## NIOS Appliances as NTP Servers

After you enable NTP on a Grid, the Grid members—including the Grid Master—can function as NTP servers to clients in different segments of the network. Similarly, after you enable NTP on an independent appliance or an HA pair, and it synchronizes its time with an NTP server, you can configure it to function as an NTP server as well. When you configure DNS anycast addressing on a Grid member and use it as an NTP server, the member can answer NTP requests from other NTP clients through the anycast IP address.

### *Grid Members as NTP Servers*



To configure a NIOS appliance as an NTP server, perform the following tasks:

- Enable the appliance as an NTP server.
- Enable authentication between the appliance and its NTP clients.
- Optionally, specify which clients can access the NTP service of the appliance.
- Optionally, specify which clients can use `ntpq` to query the appliance.

### Configuring a NIOS Appliance as an NTP Server

You can configure a Grid member—including the Grid Master—or an independent appliance or HA pair to function as an NTP server. When you enable a NIOS appliance to function as an NTP server, you can enable authentication between a NIOS appliance functioning as an NTP server and its NTP clients. When you enable authentication, you must specify the keys that the appliance and its clients must use for authentication. In a Grid, you can enter NTP authentication keys at the Grid level so that all the members can use them to authenticate their clients. You can also enter keys at the member level, if you want that member to use different keys from those set at the Grid level. After you enter the keys, you can download the key file and distribute the file to the NTP clients.

On an HA member, the NTP service runs on the active node. If there is an HA failover, the NTP service is automatically launched after the passive node becomes active and the NTP traffic uses the HA port on one of the nodes from an HA pair, instead of the LAN1 port. You might receive an error message indicating that the NTP is out of synchronization.



During another HA failover, the currently passive node becomes active again and the NTP traffic uses the LAN1 port, and the NTP is back in synchronization. For information, see [About HA Pairs](#).

To enable an appliance as an NTP server and authenticate NTP traffic between a NIOS appliance and an NTP client, perform the following tasks:

- Enable an appliance as an NTP server and define authentication keys. For information, see [Enabling an Appliance as an NTP Server](#) below.
- Optionally, define NTP access control, including KoD packet configuration. For information, see [Defining NTP Access Control](#) below.
- Optionally, configure anycast addressing for DNS and use the anycast IP address for NTP requests. For information about how to configure DNS anycast, see [Configuring Anycast Addresses](#).

### Enabling an Appliance as an NTP Server

To enable an appliance as an NTP server and add authentication keys:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox.
2. Expand the Toolbar and click **NTP** -> **NTP Member Config**.
3. In the **General** tab of the *Member NTP Properties* editor, do the following:
  - **Enable the NTP Server on this Member:** Select this option to configure a Grid Master or a Grid member as an NTP server. If you have configured DNS anycast on the appliance, it can answer NTP requests through the anycast IP address.
  - Click **Override** in the NTP Keys section to enter NTP authentication keys at the member level. The member uses these keys when acting as an NTP server and authenticates requests from NTP clients. Clear the checkbox to use the Grid-level authentication keys.
4. Click **Add** in the NTP Keys section. For information, see [Adding NTP Authentication Keys](#) above.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

After you enter the authentication keys, you can download the key file (usually called ntp.keys) and distribute it to NTP clients as follows:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab.  
**Member:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox.
2. Expand the Toolbar and click **NTP** -> **Download NTP Keys**.
3. In the *Opening ntp.keys* dialog box, save the file, and then click **OK**.
4. Distribute this to the NTP clients using a secure transport.

### Defining NTP Access Control

The NTP access control list specifies which clients can use a NIOS appliance as an NTP server. If you do not configure access control, then the NIOS appliance allows access to all clients. You can configure access control at the NTP Grid level and override that at the member level.

In addition, the NIOS appliance can accept queries from clients using ntpq, the standard utility program used to query NTP servers about their status and operational parameters. You can specify from which clients the NIOS appliance is allowed to accept ntpq queries. The appliance does not accept ntpq queries from any client, by default.

You can use an existing named ACL (access control list) or multiple ACEs (access control entries) to control which clients can use the NIOS appliance as an NTP server, as well as those clients from which the appliance can accept queries using ntpq. For information about access control, see [Configuring Access Control](#).

To specify which clients can access the NTP service of a NIOS appliance and from which clients a NIOS appliance can accept ntpq queries, and to enable or disable KoD, complete the following:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **NTP** -> **NTP Grid Config**.  
**Member:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox. Expand the Toolbar and click **NTP** -> **NTP Member Config**.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the **Access Control** tab of the *Grid* or *Member NTP Properties* editor, select one of the following to configure NTP access control:
  - **None:** Select this if you do not want to configure access control for NTP service. When you select **None**, the appliance allows all clients to access the NTP service. This is selected by default.

- **Use Named ACL for Time only:** Select this and click **Select Named ACL** to select a named ACL that contains only IPv4 addresses and networks. NTP queries do not support TSIG key based ACEs. When you select this, the appliance allows clients that have the **Allow** permission in the named ACL to use its NTP service. NTP queries from the named ACL entries specified here are denied. You can click **Clear** to remove the selected named ACL and the appliance accepts ntpq queries from those NTP clients.
- **Use Named ACL for Time + NTP Control (NTPQ):** Select this and click **Select Named ACL** to select a named ACL that contains only IPv4 addresses and networks. NTP queries do not support TSIG key based ACEs. When you select this, the appliance allows clients that have the **Allow** permission in the named ACL to use its NTP service, and for the appliance to accept ntpq queries from those clients as well. You can click **Clear** to remove the selected named ACL.
- **Use this set of ACEs:** Select this to configure individual ACEs. Click the *Add* icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows:
  - **IPv4 Address:** Select this to add an IPv4 address. Click the **Value** field and enter the IPv4 address. The default permission is **Allow**, which means that the appliance allows access to and from this IPv4 client. You cannot change the default permission. In the **Service** field, select **Time only** to allow this client for using the NTP service on the appliance; or select **Time + NTP Control (NTPQ)** to also accept ntpq queries from this client.
  - **IPv4 Network:** Select this to add an IPv4 network. Click the **Value** field and enter the IPv4 network. The default permission is **Allow**, which means that the appliance allows access to and from this IPv4 network. You cannot change the default permission. In the **Service** field, select **Time only** to allow this network for using the NTP service on the appliance; or select **Time + NTP Control (NTPQ)** to also accept ntpq queries from this network.
  - **IPv6 Address:** Select this to add an IPv6 address. Click the **Value** field and enter the IPv6 address. The default permission is **Allow**, which means that the appliance allows access to and from this IPv6 client. You cannot change the default permission. In the **Service** field, select **Time only** to allow this client for using the NTP service on the appliance; or select **Time + NTP Control (NTPQ)** to also accept ntpq queries from this client.
  - **IPv6 Network:** Select this to add an IPv6 network. Click the **Value** field and enter the IPv6 network. The default permission is **Allow**, which means that the appliance allows access to and from this IPv6 network. You cannot change the default permission. In the **Service** field, select **Time only** to allow this network for using the NTP service on the appliance; or select **Time + NTP Control (NTPQ)** to also accept ntpq queries from this network.
  - **Any Address/Network:** Select this to allow access to all IPv4 and IPv6 addresses and networks. The default permission is **Allow**, which means that the appliance allows access to and from all IPv4 and IPv6 clients. You cannot change the default permission. In the **Service** field, select **Time only** to allow clients for using the NTP service on the appliance; or select **Time + NTP Control (NTPQ)** to also accept ntpq queries from all clients.

After you have added access control entries, you can do the following:

- Select the ACEs that you want to consolidate and put into a new named ACL. Click the **Create new named ACL** icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
  - Reorder the list of ACEs using the up and down arrows next to the table.
  - Select an ACE and click the **Edit** icon to modify the entry.
  - Select an ACE and click the **Delete** icon to delete the entry. You can select multiple ACEs for deletion.
  - **Enable KoD:** When you select this checkbox, the appliance (when acting as an NTP server) sends a KoD (Kiss-o'-Death) packet to the NTP client if the client has exceeded the rate limit. The KoD packet contains the stratum field set to zero and the ASCII string in the Reference Source Identifier field set to RATE, indicating the packets sent by the client have been dropped by the server. When you clear the checkbox, the NTP server drops the packets but does not send any KoD packet to the client. This checkbox is deselected by default. For more information about KoD, see *Enabling Kiss-o'-Death for NTP* below.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Enabling Kiss-o'-Death for NTP

When an NTP server denies service to an NTP client, which has exceeded the rate limit, it typically drops the packets without notifying the client. In this case, the client, unaware of the situation, continues to transmit packets. To notify the client so it either slows down or stops the packet transmission, you can enable the NIOS appliance (when acting as an NTP server) to transmit a KoD (Kiss-o'-Death) packet. This packet contains the stratum field which is set to zero, implying the sent packet was invalid, and the ASCII string that contains RATE in the reference identifier field, indicating the status of the transmitted packet and access control. When the client receives the KoD packet, it may reduce transmission rate or stop packet transmission to the server. For more information about KoD, refer to [RFC 5905 \(Network Time Protocol Version 4: Protocol and Algorithms Specification\)](#). You can enable KoD at the Grid level and override it at the member level. For more information about enabling KoD, see [Defining NTP Access Control](#) above.

## Defining NTP Orphan Mode

The NTP orphan mode allows you to configure a stratum value that enables the Grid members to continue serving NTP uninterrupted using the disconnected NTP service in the absence of external NTP servers of the Grid. When the external NTP servers are reachable again, the Grid connects with the server to serve NTP and derive the NTP stratum values and the Grid automatically switches to the connected mode. For information about orphan mode, see [Configuring the Orphan Mode](#).

To configure a user specified stratum for NTP service on a Grid to use the disconnected NTP services, complete the following:



### Note

- Unless a special configuration is required, use the default values. In case you configure the values, keep the configuration as simple as possible.
- When you select the **Use Default** option for the stratum value of either the Grid Manager or the Grid member, you will not be able to add or edit the stratum values.
- You can use the `set ntp_stratum` CLI command in maintenance mode to set the local NTP stratum value for both the Grid Manager and member.

1. From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **NTP -> NTP Grid Config**.
2. In the **Orphan Mode** tab of the *Grid NTP* editor, specify the following:
  - **Use Default:** Select this if you want to use the default values and do not want to configure the stratum for the Grid Manager and the Grid members. When you select the **Use Default** option, NIOS uses the default stratum value of 12 for the Grid Manager and 14 for the Grid level members.
  - **Grid Manager local NTP stratum:** Specify the NTP stratum value for the Grid Manager, enter a stratum value that is between 2 to 14, both inclusive.
  - **Member local NTP stratum:** Specify the NTP stratum for the Grid members at Grid level, enter a stratum value that is between 3 to 15, both inclusive. However, the minimum NTP stratum value of the member needs to be 1 more than the Grid Manager NTP stratum value.
3. Click **Save & Close** to complete the configuration.

To configure a user specified stratum for NTP service on a member to use the disconnected NTP services, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox. Expand the Toolbar and click **NTP -> NTP Member Config**.  
To override a Grid level inherited stratum value, click **Override**. To inherit the same values as the Grid, click **Inherit**.
2. In the **Orphan Mode** tab of the *Member NTP Properties* editor, specify the following:
  - **Use Default:** Select this if you want to use the default values and do not want to configure the stratum value for the Grid Member. When you select the **Use Default** option, NIOS uses the default stratum value of 14 for the Grid member.

- **Member local NTP stratum:** Specify the NTP stratum value for the Grid member, enter a stratum value that is between 2 to 15, both inclusive. By default, if the Grid Manager is set as the NTP server for the Grid member, then the Grid member will generate the stratum value by using the stratum value of the Grid Manager.





3. Click **Save & Close** to complete the configuration.

The NTP service will restart automatically once the configuration has been modified.

## Monitoring NTP

When you enable the Grid to synchronize its time with external NTP servers, you can monitor the status of the NTP service by checking the NTP status icons in the *Member Services* panel. To access the panel, from the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then select the Manage Member Services icon in the table toolbar of the Members tab.

The following are descriptions of the NTP status icons in the *Members Services* panel. The type of information that can appear in the **Description** column corresponds to the SNMP trap messages. For information about the Infoblox SNMP traps, see [Configuring SNMP](#).

Icon	Color	Meaning
	Green	The NTP service is enabled and running properly.
	Yellow	The NTP service is enabled, and the appliance is synchronizing its time.
	Red	The NTP service is enabled, but it is not running properly or is out of synchronization.
	Gray	The NTP service is disabled.

After you upgrade the Grid to 6.6.x or later, the color of the Grid status icon changes based on the following:

- If you activate an external synchronization, or start the NTP service using the Grid Manager, or do not configure any external NTP servers, except local, then the NTP behavior remains the same and the NIOS appliance displays the Grid status icon in green.
- If you activate an external synchronization and configure one or more external NTP servers, or if the servers are in synchronization with the Grid Master, then the Grid status icon is as follows:
  - Green: NTP is synchronizing with an external server.
  - Red: NTP is synchronizing with the local server and none of the external NTP servers are reachable. This status icon also indicates if there are problems with the NTP service.
  - Yellow: NTP is synchronizing with the local server and at least one external NTP server is reachable; but there could be problems on the external server, such as an exceeded root distance error.

## Managing Extensible Attributes

Extensible attributes are identifiers that you use to further define and track a NIOS object. For example, to specify the location of a network, you can add the predefined attribute **Site** and enter a specific location for the network. You can also specify whether an extensible attribute is required for an object or restrict the values that can be entered when you create a new object.

You can also specify if an extensible attribute is inheritable by other objects in an inheritance chain. When you enable the inheritance of an extensible attribute, all descendants in the inheritance chain can inherit the extensible attribute so you do not have to configure it at all object levels. For example, if you define an extensible attribute for a network, the attribute and its value can be automatically added for DHCP ranges and fixed addresses in the network.

An extensible attribute is inheritable by descendants in an inheritance chain if its definition does not restrict it to objects that are not part of an inheritance chain. The appliance supports this inheritance chain: Network View -> Network Container -> Network -> Range -> Host/Fixed Address/Reservation. A parent object can have descendants at one or more levels. For example, a network view, network container, network or DHCP range can be a parent object and have

descendants at one or more levels, while a host, fixed address, and reservation can only be a descendant, not a parent. You can set an extensible attribute to be inheritable by selecting the **Enable Inheritance** option when you define an attribute. For more information, see [Configuring Inheritable Extensible Attributes](#) below.



#### Note

Only superusers can configure extensible attributes.

You can use predefined extensible attributes or specify new ones for different objects. For more information on supported object types and their corresponding fields for the extensible attributes, see [Subscriber Record](#). The appliance provides the following predefined extensible attributes that you can customize:

- Region
- Country
- State
- Site
- Building
- VLAN
- ReportingSite (Report Clustering)
- Parental-Control-Policy
- Proxy-All
- User-Name
- White-List
- Black-List
- Subscriber-Secure-Policy
- PC-Category-Policy



#### Note

Using the **CSV import** option, if you import DHCP network, fixed address or reservation address with Parental control extensible attributes, then the subscriber records are not created.

When you use a predefined attribute, you can modify it and change its name, but you cannot change the type of data it accepts. You can also delete predefined attributes that you do not use. All predefined attributes accept text strings. You can define other settings though, as described in [Modifying Extensible Attributes](#) below. You can also create your own extensible attributes, as described in [Adding Extensible Attributes](#) below.

For example, you can configure the predefined attribute **Site** for fixed addresses and hosts, and define a new attribute **Department** for admin groups. If you enable the option **Allow NATed Subscribers only** in subscriber site, then the extensible attributes used as the default policy for the network, fixed, and reservation address for the DHCP server members, which belong to a Subscriber Site is not supported. For more information on enabling the **Allow NATed Subscribers only** option, see, [Adding Subscriber Sites](#).

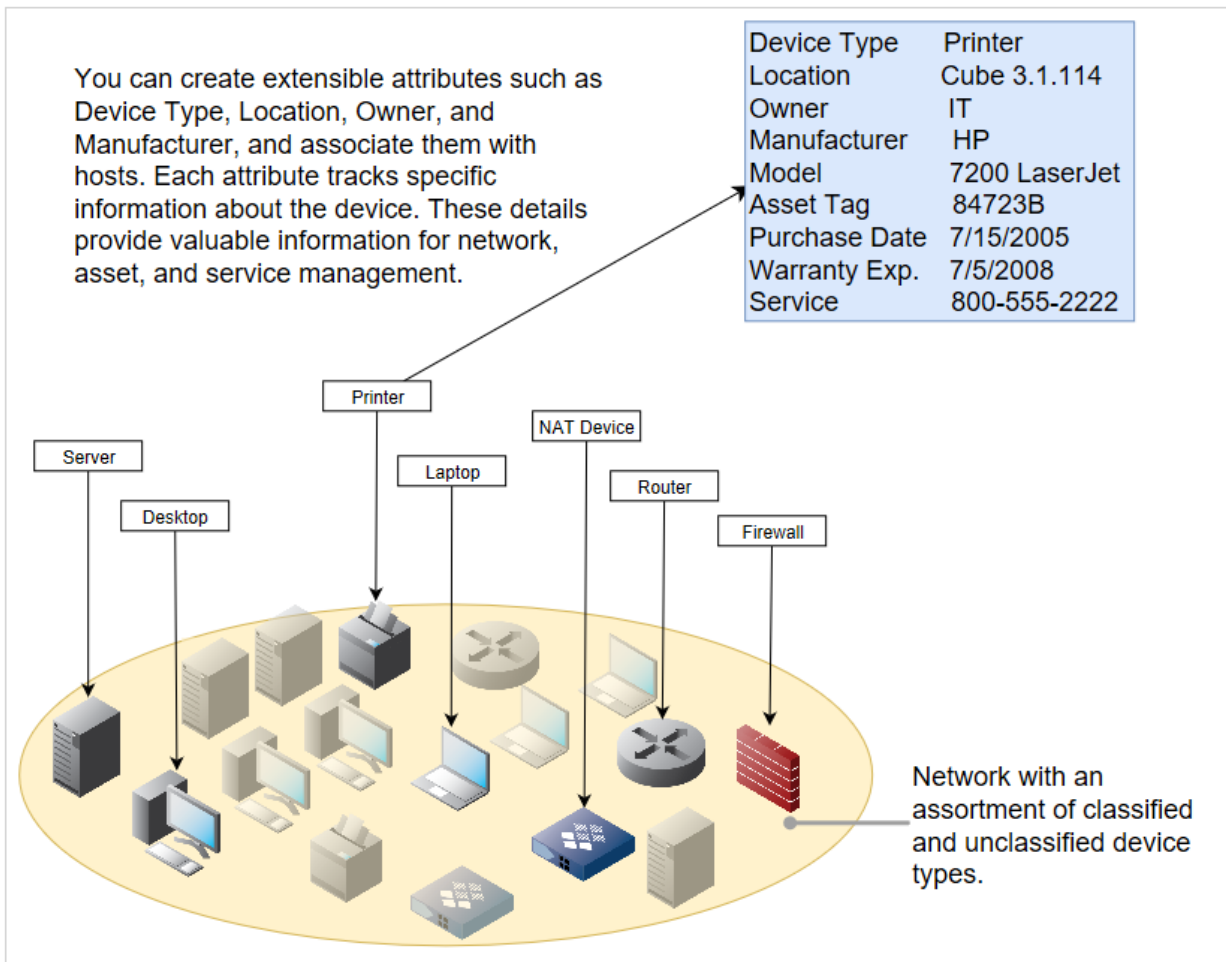
When you configure an extensible attribute, you can specify the following:

- The type of data that admins enter, such as text strings, integers, or email addresses. You can also restrict admins to a list of values.
- Whether admins can enter multiple values
- A default value
- Whether the attribute is required  
Note that You can specify whether an extensible attribute is required or not only by selecting the **Required** checkbox on the **Administration > Extensible Attributes** tab in Grid Manager. You cannot specify whether an extensible attribute is required or not by using CSV or WAPI. For more information about extensible attribute options, see [Adding Extensible Attributes](#) below.
- Whether the attribute is inheritable
- The objects associated with the attribute, such as admin groups, DNS views, or DHCP networks.
- Whether the appliance makes an entry in the audit log each time an object with the attribute is added or modified.

Activities such as additions, modifications, and deletions of inheritable extensible attributes, are recorded in the audit log. For more information about how to use the audit log, see [Using the Audit Log](#).

Using Extensible Attributes to Define Network Devices figure illustrates a network with different device types. Each device is represented as a host in the NIOS appliance database. You can configure **Device Type**, **Location** and **Owner** as required attributes for hosts. Then when admins add hosts, they will be required to enter values for these attributes in the **Extensible Attributes** tab of the *Add Host* wizard.

#### Using Extensible Attributes to Define Network Devices



After you configure extensible attributes for an object, the attributes become available in the **Extensible Attributes** tab of the wizard and editor of the corresponding object. Users then add or edit the attribute values, based on your configuration. Users can also specify attributes when searching for data and add attributes as columns in the tables of Grid Manager. For example, you can add the predefined **Site** attribute as a column in the Records panel of the **Zones** tab. For more information about adding columns to tables and customizing Tables, see [About the Grid Manager Interface](#).

Users can also group objects in smart folders according to their attributes. For example, a user can create a smart folder that contains all networks in a certain site.

Users can enable the appliance to group members by extensible attributes. For more information, see [Grouping Members by Extensible Attributes](#).

When you first enable Cloud Network Automation, NIOS installs a set of extensible attributes that are specific for cloud

usage. You should avoid adding these extensible attributes manually before enabling Cloud Network Automation. For more information, see [Extensible Attributes for Cloud Objects](#).

## Adding Extensible Attributes

To add a new extensible attribute:

1. From the **Administration** tab, select the **Extensible Attributes** tab.
2. Click the Add icon on any of the toolbars.
3. In the *Add Extensible Attribute* wizard, complete the following:
  - **Name:** Enter the name of the attribute. This is a required field and is case-sensitive. You can enter up to 128 UTF-8 characters.
  - **Type:** Specify the type of data that you want to capture for an object. Select one of the following:
    - **String:** Select this when the attribute is used to define string values, such as names. When you select this type, the wizard displays the **Number of Characters** field where you can enter the minimum and maximum number of characters that users can enter.
    - **List:** Select this when you want to define a list of values for the attribute. Users can then select a value from this list. For example, if you want to restrict an attribute to five specific values, you can define the attribute as a **List** and then list the five values in the *List Values* section. When a user uses the attribute, they are limited to selecting from one of the five values. When you select **List**, the wizard displays the List of values table, where you add the allowed values. These values appear in the drop-down list when a user defines the attribute. Click the Add icon to enter values in the table. You can enter up to 64 UTF-8 characters for each value. You can also modify list values at a later time. When you modify list values, all object attributes using the modified values are updated to the new values. You can also delete values from the list. Note that when you delete a list value, all attributes using the deleted values are removed from the objects. For objects with multiple attribute values, only the deleted values are removed. You can also move a value up or down in the list.
    - **Integer:** Select this when the attribute is used to track whole numbers, such as serial numbers. When you select this type, the wizard displays the **Value Limits** fields where you can enter the range of allowed values. Note that you cannot change your entries in the **Value Limits** fields if you modify the attribute at a later date.
    - **Email:** Select this when the attribute is used for email addresses. Email addresses are entered in the format *user@domain.com*.
    - **URL:** Select this when the attribute is used for tracking URLs (Uniform Resource Locators). URLs must be entered in a valid format.
    - **Date:** Select this when the attribute is used for dates. The date value is in YYYY-MM-DD format.
  - **Comment:** Enter additional information about the attribute. You can enter up to 256 UTF-8 characters.
4. Click **Next** and complete the following:
  - **Enable Inheritance:** Select this checkbox if you want to allow the extensible attribute and its values to be inherited by descendants in an inheritance chain. When you select this checkbox, inheritance is enabled for network related objects only. When you select this checkbox and restrict an attribute to certain objects, then the extensible attribute and its value will be inherited by those objects only. Note the following:
    - If you create an extensible attribute with inheritance disabled and later enable it, the *Descendant Actions* dialog box may be displayed with the available options for adding an extensible attribute. For more information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#) below.
    - If you create an extensible attribute with inheritance enabled and later disable it, the *Descendant Actions* dialog box may be displayed with the available options for deleting an extensible attribute. For more information, see [Deleting Inheritable Extensible Attributes Associated with Parent Objects](#) below.
  - **Allow multiple values:** Select this checkbox if you want to allow multiple values for this attribute to be set on an object. You cannot change this value for predefined attributes. Once you select this checkbox, you cannot deselect it. That is, you cannot allow a single attribute to be set if you selected to allow multiple values to be set on an object.

### Restricting synchronization of extensible attributes



- **Disable sync to MGM:** Select this checkbox to disable synchronization of extensible attributes from the managed Grid to the Multi-Grid Master. This checkbox is available only on the managed Grid when it remains joined with the Multi-Grid Master.
  - **Default Value:** Enter the default value that the appliance displays for the attribute. Leave this blank if there is no default value for this attribute. If the attribute type is **String**, you can enter up to 256 UTF-8 characters. If the attribute type is **List**, the value must be one of the list values and can be up to 64 UTF-8 characters.
  - **Required:** If you select this option, it is required to enter a value for this attribute when adding or modifying the corresponding object in the GUI.
  - **Recommended:** If you select this option, it is recommended to enter a value for this attribute when adding or modifying the corresponding object in the GUI.
  - **Optional:** This is selected by default. By selecting this option, you may or may not enter a value for this attribute when adding or modifying the corresponding object in the GUI.
  - **Restrict to Specific Object Types:** Click the Add icon to select the object type with which you want to associate the attribute. The appliance adds a row to the table. To delete an object type, select an object type and click the Delete icon. By default, the appliance associates an extensible attribute with all the supported object types.
  - **Log Attribute Values When Objects are Updated:** Select this checkbox if you want the appliance to make an entry in the audit log each time an object with this attribute is added or modified. When you select attribute values for audit, they are included in all the audit log entries. For more information about the audit log, see [Using the Audit Log](#).
  - **Allow cloud members to have the following access to this extensible attribute:** Select this if you are configuring this extensible attribute for Cloud Network Automation. When you select this checkbox, the Cloud Platform Appliance can access this extensible attribute and perform requested tasks based on the cloud API requests. This function is enabled by default for all cloud specific extensible attributes. Note that if you disable this function for any cloud attributes, you will receive an error when you try to perform tasks that involve these attributes through cloud API requests. You can select **Read/Write** or **Read only**. For more information about this feature, see [Deploying Cloud Network Automation](#).
    - **Read/Write (and disallow Write access from the GUI and the standard API):** When you select this, the Cloud Platform Appliance can access and modify the value of this extensible attribute based on the tasks requested only through cloud API requests. You cannot modify this attribute using Grid Manager or the Infoblox API.
    - **Read only:** When you select this, the Cloud Platform Appliance can access this extensible attribute and report the value based on the cloud API requests, but it cannot modify the value. You receive an error if you try to modify this attribute when this option is selected.
5. Save the configuration and click **Restart** if it appears at the top of the screen.



#### Warning

The Cloud Platform member is evicted from the Grid Master if you modify the extensible attribute using Grid Manager.

Grid Manager adds the attribute to the **Extensible Attributes** tab in either the wizard or editor for the specified object types.



#### Note

Infoblox recommends that you define values for mandatory extensible attributes using the Grid only and do not use PAPI or RESTful API to define values.

## Configuring Inheritable Extensible Attributes

An extensible attribute can be inherited by descendants when it is at the top or in the middle of the inheritance chain. When you add a new extensible attribute to a parent object, the same extensible attribute may or may not already exist at the descendant levels. If the extensible attribute exists on a descendant, you can choose to have the descendant inherit the value from the parent, or retain the original value from the descendant. When the extensible attribute does not exist on the descendant, you can choose to have the descendant either inherit the extensible attribute and its value from the



parent or not inherit anything from the parent.

When you add a range, host, fixed address or IPv4 reservation to a parent object which has inheritable extensible attributes, the newly added object can inherit extensible attributes from the parent object. For example, if you create an IPv4 network with inheritable extensible attributes, and then add a host, the values you specified for the extensible attributes while creating the network can be inherited by the host.

To assist you in identifying whether an extensible attribute value is inherited or overridden, the appliance displays the inheritance state of an attribute in the **Inheritance State** column of an extensible attribute. This column is displayed only for objects that support inheritance. For more information about how to view inheritance states, see [Modifying Inheritable Extensible Attributes](#) below.

Following are the supported inheritance states:

#### *Inheritance States*

Inheritance	State Description
Inherited	The extensible attribute inherits its value from the parent. You cannot edit the value of an attribute when the inheritance state is set to <b>Inherited</b> . You can change the state to <b>Overridden</b> and then change the value of the attribute or change the state to <b>Not Inherited</b> to remove the inherited value.
Overridden	The extensible attribute overrides the value inherited from the parent. You can change the state to <b>Inherited</b> and restore the original inherited value or change the state to <b>Not Inherited</b> and remove the inherited value.
Not Inherited	The extensible attribute can inherit its value from the parent, but the attribute does not exist on the descendant. You can change the state to <b>Inherited</b> and restore the original inherited value or change the state to <b>Overridden</b> and change the value of the attribute. Note that when the state of an inheritable extensible attribute is <b>Not Inherited</b> , the corresponding attribute will not be added as a new extensible attribute for objects that are currently not inheriting this extensible attribute.
No Parent	The inheritance state is set to <b>No Parent</b> when an object has a parent, but the parent does not have an extensible attribute or the parent's extensible attribute is set to <b>Not Inherited</b> .
Disabled	Extensible attribute inheritance is not enabled for the attribute.
No Change	The extensible attributes of the selected objects do not have the same inheritance state for all objects. This state allows you to retain the current state on the selected objects. Note that this state is only seen in the <i>Multi-object Extensible Attributes</i> editor.

When you add an inheritable extensible attribute to an object, if there are descendants of this object the *Descendant Actions* dialog box is displayed which will provide options for the descendants. Following is a summary of these options:

- Retain the original value of the attribute for all descendants.
- Inherit the extensible attribute and its value from the parent object.
- Inherit the extensible attribute and its value when it does not exist on descendants.
- If the extensible attribute does not exist on the descendant, do not add it.
- If you are deleting an inherited extensible attribute from a parent object you can retain or remove the extensible attribute from the object's descendants.

To configure default descendant actions for inheritable extensible attributes:

1. From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.
2. In the **Extensible Attribute Inheritance** tab, complete the following:

**When adding an extensible attribute that already exists on a descendant:**

- **Keep the descendant's existing value and change the inheritance state to Override:** Select this if you want to retain the existing extensible attribute values for all direct descendants, irrespective of the values you define at the parent level. The inheritance state for all direct descendants will be set to **Overridden**. Note that this is applicable only when you add a new extensible attribute to the parent object that already exists on the descendant. If you modify the value of an existing extensible attribute that is already inherited by the descendant, and select the above option in the Descendant Actions dialog box, then the new value will be inherited by the descendant, but the inheritance state will remain **Inherited**. For example, consider a network 10.0.0.0/16 that has an extensible attribute **Site** with the value **SA** (native). When you add another network 10.0.0.0/24, extensible attribute **Site** inherits its value, **SA**, from the parent object. Now, if

you add network 10.0.0.0/8, assign extensible attribute **Site** and set its value to **NY**, then when you choose this option, the value of **Site** will remain as **SA**, but the inheritance state will be changed to **Overridden** for network 10.0.0.0/16; however, network 10.0.0.0/24 will still have its value as **SA** for **Site** with the inheritance state set to **Inherited**.

- **Inherit the parent's value and change the inheritance state to Inherit:** Select this to inherit the extensible attribute values from the parent for all descendants. The inheritance state for all descendants will be set to **Inherited**.
- **Change the inheritance state to Inherit only if the descendant's value is the same as the parent's value. Otherwise, change the state to Override:** Select this to set the inheritance state to **Inherit** if the descendants have the same extensible attribute value as the parent. Otherwise, retain the original extensible attribute value on the descendants and change the inheritance state to **Overridden**.

**When adding an extensible attribute that does not exist on a descendant: Do not inherit the value from the parent and change the inheritance state to Not Inherited:** Select this if the extensible attributes do not exist on the descendants and you do not want them to inherit the attributes from the parent. The inheritance state is set to **Not inherited**.

- **Inherit the value from the parent and change the inheritance state to Inherited:** Select this if you want all descendants to inherit extensible attributes from the parent, and the inheritance state for all descendants will be set to **Inherited**.

**When deleting an extensible attribute:**

- **Keep the descendant's value and change the inheritance state to No Parent:** Select this if you want to preserve extensible attributes on all descendants when you delete an inheritable extensible attribute. The inheritance state for direct descendants will be set to **No Parent**.
- **If the current inheritance state is Inherited, remove the extensible attribute. If the current inheritance state is Overridden, keep the value and change the inheritance state to No Parent:** Select this if you want to remove the extensible attributes that are inherited by the descendants. If the inheritance state of the extensible attributes is set to **Inherited** on the descendant, the attributes will be removed; however, if the inheritance state is set to **Overridden**, then the state will be changed to **No Parent**.

3. Save the configuration.

For more information about how to configure inheritable extensible attributes, see [Configuration Examples for Inheritable Extensible Attributes](#).

### Admin Permissions and Inheritable Extensible Attributes

Permissions for descendant objects can affect the results of the actions that are chosen in the *Descendant Actions* dialog box:

- When you add an extensible attribute to the parent object: The descendants to which you have read-write permission will behave as expected with any of the chosen options in the *Descendant Actions* dialog box.
- When you change the extensible attribute value on the parent object: The descendants that have the same extensible attribute set to **Inherited** will be automatically changed to the new value, even though you may not have write permission for those descendants.
- When you select to preserve descendant values while removing an extensible attribute associated with the parent object, values will be preserved even if you do not have write read-write permission for those descendants.
- When you select to remove an extensible attribute on descendants when removing a parent's extensible attribute, an error message will be displayed if you do not have read-write permission to some of the descendants.

### Guidelines for Configuring Inheritable Extensible Attributes

- When you add an inheritable extensible attribute to a parent object, you can choose to have descendants inherit or override the parent's extensible attribute value. You can also choose that the extensible attribute not be added to a descendant.
- When you add a new parent with an inheritable extensible attribute, the options for changes to descendants remain the same as when you add a new inheritable extensible attribute to a parent. For more information, see [Configuring Inheritable Extensible Attributes](#) above.
- When you add a new descendant to the existing parent with inheritable attributes, the descendant inherits all the extensible attributes. However, you can select if you want to inherit or override the values. If you set the inheritance state to **Not Inherited**, then the extensible attribute will not exist on the descendant, but you can later change the state to **Inherited** or **Overridden**. For more information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#) below.

- When you delete an inheritable extensible attribute associated with the parent, you can either preserve the extensible attribute values on the descendants or delete the inherited extensible attributes. For more information, see *Deleting Inheritable Extensible Attributes Associated with Parent Objects* below.
- When you delete a parent object and if there is grandparent, then the extensible attribute will be re-parented when you choose preserve. The current inheritance state of the attribute will be retained. If you delete a parent object and if there is no grandparent, then the inheritance state of the extensible attribute is changed to **No Parent** when you choose preserve.
- When you split a network, the extensible attribute will be copied to the newly created networks. For inheritable extensible attributes, the newly created network inherits the extensible attributes and the state is set to **Inherited**. For more information, see *Managing Inheritable Extensible Attributes at the Parent and Descendant Level* below.
- When you join two networks to form a larger network, the *Descendant Actions* dialog box is displayed with the following options:

**When joining networks, select the action(s) you want to apply to the merged networks:**

- **Preserve extensible attributes for all descendants of the merged networks and change the inheritance state to No Parent:** Select this if you want to preserve the extensible attributes for all descendants of the merged networks. The inheritance state of the attributes will be changed to **No Parent**.
- **Remove extensible attributes from descendants of the merged networks:** Select this if you want to remove extensible attributes that are inherited by descendants.  
Note that the options above apply only to extensible attributes which no longer have a parent, due to the merge. If the extensible attributes on descendants are also on the resulting merged network, then they will retain their current state.  
When you add multiple inheritable networks, new networks will automatically inherit all extensible attributes from the parent.

## Managing Inheritable Extensible Attributes at the Parent and Descendant Level

You can define if descendants will inherit values from the parent when a new extensible attribute is added to the parent. You can also choose to override the values of the extensible attributes on the descendants.

When you delete an existing attribute, you can choose to either preserve the values at the descendant level or delete the values inherited by the descendants.



### Note

The *Descendant Actions* dialog box is displayed only when an object has descendants and you are modifying extensible attributes that affects those descendants. However, the dialog box is always displayed when a join is performed for a network that has inheritable extensible attributes.

The following section describes configuration changes for inheritable extensible attributes:

1. **Network Container:** From the **Dashboards** tab, select the **Tasks** tab -> click **Add Networks**. Select a network, enter the required details. You can edit the inheritable extensible attributes that are displayed automatically. If this is a parent object, then you can add extensible attributes.
  - IPv4 Network:** From the **Data Management** tab -> select the **DHCP** tab -> **Networks** tab. In the **Networks** section, select **IPv4 Network** from the Add drop-down menu. In the *Add IPv4 Network* wizard, enter the attributes in the **Extensible Attributes** tab after specifying the required details.
  - IPv6 Network:** From the **Data Management** tab -> select the **DHCP** tab -> **Networks** tab. In the **Networks** section, select **IPv6 Network** from the Add drop-down menu. In the *Add IPv6 Network* wizard, enter the attributes in the **Extensible Attributes** tab after specifying the required details.
  - IPv4 Range:** From the **Data Management** tab > select the **DHCP** tab -> **Networks** tab -> **Networks** tab -> network > click **addr\_range**, select **Range** from the Add drop-down menu. In the *Add IPv4 Range* wizard, enter the attributes in the **Extensible Attributes** tab after specifying the required details.
  - IPv6 Range:** From the **Data Management** tab > select the **DHCP** tab -> **Networks** tab -> **Networks** tab -> network > click **addr\_range**, select **Range** from the Add drop-down menu. In the *Add IPv6 Range* wizard, enter the attributes in the **Extensible Attributes** tab after specifying the required details.
  - Zones:** From the **Data Management** tab > select the **DNS** tab -> **Zones** tab -> click **Add**. In the *Add Zone* wizard, enter the attributes in the **Extensible Attributes** tab after specifying the required details.
  - DNS View:** From the **Data Management** tab > select the **DNS** tab -> In the toolbar select **Add** -> **Select DNS View**. In the *Add DSN View* wizard, enter the attributes in the **Extensible Attributes** tab after specifying the required details.

**Host:** From the **Data Management** tab > select the **DNS** tab - > **Zones** tab - > click Add. In the Add section, select **Host** or **Bulk Host** from the Add drop-down menu. In the *Add DNS View* wizard, enter the attributes in the **Extensible Attributes** tab after specifying the required details.

**Record:** From the **Data Management** tab > select the **DNS** tab - > **Zones** tab - > click Add. In the Add section, select **Record** from the Add drop-down menu select the record type. In the *Add DNS View* wizard, enter the attributes in the **Extensible Attributes** tab after specifying the required details.

**VLAN Range:** From the **Data Management** tab > select the **VLANS** tab - > click Add, select **VLAN Range** from the Add drop-down menu. In the *Add VLAN Range Wizard*, enter the attributes in the **Extensible Attributes** tab after specifying the required details.

**VLANS:** From the **Data Management** tab > select the **VLANS** tab - > click Add, select **VLAN** from the Add drop-down menu. In the *Add VLAN* wizard, enter the attributes in the **Extensible Attributes** tab after specifying the required details.

2. You can either add new extensible attributes to the parent object or modify original extensible attribute values. Click on the extensible attribute value displayed in the **Value** column of the respective attribute to modify the original value or click the Add icon to add a new attribute.
3. Select a state from the drop-down list displayed in the **Inheritance State** column. Note that you can only change the inheritance state of a descendant. You must select **Overridden** from the drop-down list to enter a new value. For more information about inheritance states, see *Inheritance States* table below. When an object has a parent and the parent does not have the object's inheritable extensible attribute, then the inheritance state of the extensible attribute is set to **No Parent** and the state cannot be changed.
4. **Select the inheritable extensible attributes for which you want to modify descendant actions:** Select this checkbox if you would like to apply the actions of the *Descendant Actions* dialog box for existing extensible attributes. Before you select this checkbox, you must select the extensible attributes which will be affected by the actions of the *Descendant Actions* dialog box.  
Note this checkbox is not displayed for hosts, fixed addresses, and reservations since they do not have descendants.
5. In the *Descendant Actions* dialog box, select options that will be applied for descendant objects as described in *Configuring Inheritable Extensible Attributes*.  
The *Descendant Actions* dialog box displays all the mentioned options when you perform add and delete operations simultaneously. Consider an example where you add a new inheritable extensible attribute **Site**, and delete an existing inheritable attribute **Region** from the parent object, and then click **Save** to save both changes. In this case, the *Descendant Actions* dialog box displays all the options.
6. Save the configuration.

## Viewing Extensible Attributes

To view the configured extensible attributes, from the **Administration** tab, select the **Extensible Attributes** tab. The panel displays the following information:

- **Name:** The name of the extensible attribute.
- **Type:** The type of data defined by the attribute.
- **Comment:** Comments entered for the extensible attribute.
- **Required:** Indicates whether users are required to complete this field.
- **Restricted to Objects:** The object types that are associated with the attribute.
- **Inheritance Enabled:** Indicates whether inheritance is enabled or not. You can do the following in this panel:
  - Sort the displayed data in ascending or descending order by column.
  - Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
  - Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
  - Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#)
- Group results by extensible attributes. For more information, see [Grouping Results by Extensible Attributes](#) below.
- Add or delete extensible attributes.
- Print or export the data.

## Modifying Extensible Attributes

When you modify an extensible attribute, all objects using the modified attributes are updated. You can perform inline editing by double-clicking the row of data that you want to modify. The appliance displays the inline editing editor in the selected row. Click **Save** after modifying the data. Note that you cannot edit extensible attributes that have multiple values.

To modify an extensible attribute:

1. In the **Administration** tab, select the **Extensible Attributes** tab.
2. Select the attribute and click the Edit icon.
3. In the **General** tab of the *Extensible Attributes* editor, you can only change the name of the attribute. You cannot change the data type. The data type for predefined attributes is string.
4. In the **Additional Properties** tab, you can modify any of the fields described in step 4 of Adding Extensible Attributes.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Modifying Inheritable Extensible Attributes

When values are inherited by a descendant, the inheritance state of the inherited extensible attribute is displayed as **Inherited**. You can select **Overridden** and specify a new value or select **Inherited** to retain the same value as the parent. If you select **Not Inherited**, the extensible attribute and its value will not be inherited. The inherited value will have a strike-through and you cannot edit the value when the state is set to **Not Inherited**.

In addition to the attribute values, the **Value** column of an inheritable extensible attribute also displays the name of the source and the object type of the extensible attribute. For example, a *Network Container* has a descendant, *Network*, which inherits an extensible attribute value from *Network Container* and *Network* has a descendant, *Fixed Address* that inherits the same extensible attribute value. In this case, *Fixed Address* shows *Network Container* as the source.

The following table displays various inheritance states and corresponding changes to source and object types that are displayed in the **Value** column of an extensible attribute.

Inheritance State	Source and Object Type in the Value Column
If an extensible attribute is a native attribute (an object which is at the top of the inheritance chain, or does not have ancestors),	<b>Source</b> is not displayed in the <b>Value</b> column. This column will not display the source details, if none of the selected objects support inheritance.
If the state of an extensible attribute is set to <b>Inherited</b> ,	then the <b>Source</b> and object is displayed. You cannot change the value of the extensible attribute.
If the state of an extensible attribute is set to <b>Overridden</b> or <b>Not Inherited</b> ,	then the <b>Source</b> will have a strike-through. You can change the state of such extensible attributes. You cannot change the value of the extensible attribute when the inheritance state is set to <b>Not Inherited</b> .

To modify the value and inheritance state of an inheritable extensible attribute:

1. **For IPv4 and IPv6 Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** → *network* -> *addr\_range* checkbox, and then click the Edit icon.  
**For IPv4 Range, IPv6 Range, Fixed Address, Reservation, and Host:** From the **Data Management** tab > select the **DHCP** tab -> **Networks** tab -> **Networks** tab -> *network* > click *addr\_range*, click the Edit icon.  
**For DNS View:** From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> select the checkbox of the **DNS View** -> click the Edit icon.  
**For VLANs:** From the **Data Management** tab, select the **VLAN** tab -> select the checkbox of the **VLAN Range / VLAN** -> click the Edit icon.
2. In the editor, click the **Extensible Attributes** tab, select the checkbox of the respective attribute.
3. At the parent level, click on the value you want to change and enter the new value.  
 At the descendant level, click on the value you want to change and enter the new value. Note that you can change the value only when the inheritance state is set to **Overridden**.



4. Select a state from the drop-down list displayed in the **Inheritance State** column. Note that you can only change the inheritance state of a descendant. You must select **Overridden** from the drop-down list to enter a new value. For more information, see [Inheritance States](#) table above.
5. **Select the inheritable extensible attributes for which you want to modify descendant actions:** Select this checkbox if you would like to apply the actions of the *Descendant Actions* dialog box for existing extensible attributes. Before you select this checkbox, select the extensible attributes which will be affected by the actions of the *Descendant Actions* dialog box. For more information about the *Descendant Actions* dialog box, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#) above.
6. Save the configuration.

## Deleting Extensible Attributes

When you delete an extensible attribute, the appliance removes the attribute. All the attribute values set on the selected object types are removed from those objects. Once deleted, the attribute no longer exists in the system. Deleted attributes are not moved to the Recycle Bin. This operation might take a long time depending on the amount of data that needs to be deleted.

To delete extensible attributes:

1. In the **Administration** tab, select the **Extensible Attributes** tab.
2. Select the attribute and click the Delete icon.
3. When the confirmation dialog box appears, click **Yes**.

## Deleting Inheritable Extensible Attributes Associated with Parent Objects

When you remove an inheritable extensible attribute, which is associated with a parent object, you can choose to retain the descendant extensible attribute or remove it from all the descendants.

Note that you cannot delete extensible attributes that have the inheritance state set to **Overridden**, **Inherited**, and **Not Inherited**. You can delete an extensible attribute that is directly assigned to the object and has its inheritance state set to **No Parent** or if the inheritance state is **Disabled**.

To remove an inheritable extensible attribute associated with a parent object:

1. **IPv4 and IPv6 Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr\_range* checkbox, and then click the Edit icon.  
**For IPv4 Range, IPv6 Range, Fixed Address, Reservation, and Host:** From the **Data Management** tab **DHCP** -> tab -> **Networks** tab -> **Networks** tab -> *network* > click *addr\_range*, select the object and click Edit icon.  
**For DNS View:** From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> select the checkbox of the **DNS View** -> click the Edit icon.  
**For VLANs:** From the **Data Management** tab, select the **VLAN** tab -> select the checkbox of the **VLAN Range / VLAN** -> click the Edit icon.
2. In the editor, click the **Extensible Attributes** tab, select the attributes and then click the Delete icon.
3. When you click **Save and Close**, the *Descendant Actions* dialog box is displayed automatically with the following options:  
**Select the action(s) you want to apply to descendant objects that have the following extensible attribute(s):**  
**When deleting an extensible attribute:**
  - **Keep the descendant's value and change the inheritance state to No Parent:** Select this if you want to preserve extensible attributes for all descendants. The inheritance state of the extensible attribute changes to **No Parent**.
  - **If the current inheritance state is Inherited, remove the extensible attribute. If the current inheritance state is Overridden, keep the value and change the inheritance state to No Parent:** Select this if you want to remove the extensible attributes that are inherited by the descendants. If the inheritance state of the extensible attributes is set to Inherited on the descendant, the attributes will be removed; however, if the inheritance state is set to Overridden, then the state will be changed to No Parent.
4. Click **Yes** to save the configuration or **No** to exit.



### Note

The deleted extensible attributes will not be moved to the Recycle Bin and you cannot restore extensible attributes that are deleted.

After a superuser admin configures the attributes of an object, they become available in the wizard and editor of the object. This section describes how users can then add and manage the attributes that were configured.

Grid Manager displays the required extensible attributes in the **Extensible Attribute** tab. You must enter values for all required attributes. If an object does not have required attributes, you can add the available optional attributes.

In the **Extensible Attribute** tab of an object, such as a network or host record, you can do the following:

- Enter values for extensible attributes
- Add attributes
- Change the inheritance state of an extensible attribute
- Select if descendants must inherit extensible attribute values from its parent
- Delete optional attributes

To enter values for the extensible attributes of an object:

1. Open the editor of the object. For example, to enter values for the attributes of a network, select it and click its **Extensible Attributes** tab.
2. Click the Value column of the attribute. You must enter values for all required attributes.
3. Depending on the required attribute type, either enter or select a value for the attribute from the Value column.
4. Based on whether the attribute is inheritable, the values are displayed in the **Inheritance State** column. This value can be set to **Inherited**, **Overridden** or **Not inherited**. If the object is at the top of the inheritance chain (Network View), then the inheritance state is not displayed. The inheritance state is set to **No Parent** only if an object has a parent, but the parent does not have the inherited extensible attribute. This column is not displayed if all selected objects do not belong to the supported inheritance chain. Example: Zones, DNS View, DNS records, etc.

To add attributes:

1. Click the Add icon. Grid Manager adds a row to the table with the default attribute displayed.
2. Click the default attribute and expand the list of available attributes.
3. Select an attribute from the drop-down list.
4. Enter or select a value for the attribute from the Value column. To delete an attribute:
5. Click the checkbox beside the attribute you want to delete.
6. Click the Delete icon.



#### Note

You cannot delete an extensible attribute which has its inheritance state set to **Inherited**, **Overridden**, and **Not Inherited**. You can delete an extensible attribute that is directly assigned to the object and has its inheritance state set to **No Parent** or if the inheritance state is **Disabled**.

To delete all attributes:

1. Click the **Attribute Name** checkbox.
2. Click the Delete icon.  
Note that you can delete only attributes that are not required. If you have one or more required attributes, you cannot use the delete all function.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

### Editing Multiple Extensible Attribute Values

You can also manage the extensible attributes of multiple objects at the same time. For example, you can select several zones, and view and modify their extensible attributes all at once in the *Multi-Select Edit Extensible Attributes* editor.

Note that Grid Manager may not apply the changes you made to all the selected objects. It applies the change to objects that meet the following criteria:

- You have read/write permission to the object.
- The selected object is not locked by another user or does not have a scheduled pending task.
- If the attribute was restricted to certain object types, the object must be one of those types. To edit multiple extensible attribute values:

1. Select the objects whose extensible attributes you want to modify. You can select specific objects or select all objects in a dataset, as described in [Selecting Objects in Tables](#).

2. Expand the Toolbar and click **Extensible Attributes**.

Grid Manager displays the *Multi-Select Edit Extensible Attributes* dialog box which lists the extensible attributes of the selected objects. It displays the following information for each attribute:

- **Attribute Name:** This field displays the name of the extensible attribute associated with the selected object.
- **Value:** If the selected objects have the same value for the attribute, Grid Manager displays that value in this field. If the selected objects have different values for the attribute or if some have values and others do not, this field displays **Multiple Values** and the cell is highlighted in gray. An attribute can have multiple rows if it allows multiple values. Grid Manager displays the values that all objects have in common, if any. Otherwise, it displays **Multiple Values**. This column displays the source for inherited extensible attributes only. Note that when you add new extensible attributes, edit values of existing extensible attributes or delete an extensible attribute, then the *Descendant Actions* dialog box is displayed, even if the objects do not have any descendants. For more information about Source values, see *Modifying Inheritable Extensible Attributes* above.
- If you select objects that have the same inherited extensible attributes, but objects have different parents, then the **Source** column will display **Multiple Ancestors**.
- If the inheritance state of an extensible attribute is **Not Inherited**, then the extensible attribute will not be added as a new extensible attribute to objects that are currently not inheriting this extensible attribute.
- **Inheritance State:** This field displays the inheritance state of an extensible attribute. The column value can be **Inherited**, **Not Inherited**, **No Parent**, **No Change** or **Overridden**. This column is not displayed if all selected objects do not belong to the supported inheritance chain. Example: Zones, DNS View, DNS records, etc. If extensible attributes for the selected objects have the same inheritance state, then the respective inheritance state is displayed in this column. When objects have different inheritance states, this column displays **No Change**, so that the current inheritance state is retained on the selected objects. If you then change the inheritance state of an extensible attribute to a specific state, the corresponding attribute will be changed to the selected inheritance state on all selected objects where the extensible attribute is currently inherited. If the object is at the top of the inheritance chain (Network View), then the inheritance state is not displayed. The inheritance state is set to **No Parent** only if an object has a parent, but the parent does not have the inherited extensible attribute. For more information about inheritance states, see *Inheritance States* table above.
- **Required:** This field displays **Yes** if the attribute is required in at least one object associated with the attribute. It displays **No** if the attribute is not required in any of the objects.

3. You can do the following:

- Change the value of an attribute. Depending on the attribute type, select the value and either enter a new value or select one from the drop-down list.
- Add an attribute to the selected objects. Click the Add icon. In the **Attribute Name** field of the new row, select an attribute from the list of available attributes and specify its value. If the attribute that you added was configured as a required attribute, the **Required** field displays **Yes**. Otherwise, it displays **No**.
- Delete an attribute. You can delete an attribute if it is not required. Select the attribute and click the Delete icon.

4. Click **OK** when you are finished modifying the extensible attributes.

Grid Manager applies your changes to the applicable objects. This operation might take a long time, depending on the amount of data being modified. You can choose to run this operation in the background, as described in [Tasks Dashboard](#).

## Grouping Results by Extensible Attributes

You can enable the appliance to group members and services with the same extensible attributes. Grid Manager displayed group data with the same value for the specified extensible attribute.

To group results by extensible attributes:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab or **Services** tab.  
From the **Data Management** tab, select the **DHCP** or **DNS** tab -> **Members/Servers** tab.
2. Complete the following to group members with the same extensible attribute value:
  - **Group Results:** Select this checkbox to enable the appliance to group members by extensible attributes.
  - **Group By:** Select an extensible attribute by which you want to group members from the drop-down list. Grid Manager displays data per group of members configured with the same extensible attribute value.



To add additional Group By filter, click the + icon, and then select a value from the drop-down list. You can apply up to 10 Group By filters. You can also delete a filter by clicking the - icon.

- **Do not inherit the value from the parent and change the inheritance state to Not Inherited:** Select this if the extensible attributes do not exist on the descendants and you do not want them to inherit the attributes from the parent. The inheritance state is set to **Not inherited**.
- **Inherit the value from the parent and change the inheritance state to Inherited:** Select this if you want all descendants to inherit extensible attributes from the parent, and the inheritance state for all descendants will be set to **Inherited**.

## Configuration Examples for Inheritable Extensible Attributes

All examples in this section are based on the inheritance chain Network View -> Network Container -> Network -> Range -> Host/Fixed Address/Reservation, in which network view is at the top level and host, fixed address and reservation at the bottom of the inheritance chain.

### Example 1

When you add an extensible attribute to the top object, the inheritance state is set to **No Parent**. For example, if you add a new inheritable extensible attribute, **Building**, to a network view, the inheritance state of this extensible attribute is set to **No Parent** for the network view.

### Example 2

When you add an extensible attribute **Site** to the parent object **Network** that has a descendant **Range**, you can define **Site** as inheritable and add it to the **Network**. The descendant, **Range**, may or may not have the same extensible attribute. Infoblox displays a list of options that lets you either inherit the value or retain or override the existing value of the extensible attribute at the descendant level. Another option is to inherit the value of **Site**, only if the value for this attribute in **Range** is same as that in **Network**. You can also decide if **Range** should acquire the same value for **Site**, if it is not defined for **Range**. This change can be inherited by the descendants of **Range**.

Depending on your configuration, the inheritance state of the extensible attribute can display **Inherited**, **Overridden** or **Not Inherited**. If the object is at the top of the inheritance chain (Network View), then the inheritance state is not displayed. The inheritance state is set to **No Parent** only if an object has a parent, but the parent does not have the inherited extensible attribute.

### Example 3

Examples in this section show different results when you add a new inheritable extensible attribute to an object located at the top or in the middle of the inheritance chain based on the following:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container			
10.0.0.0/16	Network	Region	ABC	Native
10.1.0.0/16	Network			

#### Example 3.1: Add an extensible attribute Region with value DEF to 10.0.0.0/8

You select the following options for the existing extensible attribute:

- For descendants that already have this extensible attribute, the existing extensible attribute will always be set to **Inherit**.
- For descendants that do not have this extensible attribute, the descendants will inherit this extensible attribute.

**Result:**

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	DEF	Native
10.0.0.0/16	Network	Region	DEF	Inherited from 10.0.0.0/8
10.1.0.0/16	Network	Region	DEF	Inherited from 10.0.0.0/8

**Example 3.2: Add an extensible attribute Region with value DEF to 10.0.0.0/8**

You select the following options for the existing extensible attribute:

- For descendants that already have this extensible attribute, the existing extensible attribute will always be set to **Override**.
- For descendants that do not have this extensible attribute, the descendants will not inherit this extensible attribute (extensible attribute is set to **Do not Inherit**).

**Result:**

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	DEF	Native
10.0.0.0/16	Network	Region	ABC	Overridden
10.1.0.0/16	Network	Region		

**Example 3.3: Add an extensible attribute Region with value DEF to 10.0.0.0/8 8**

You select the following options for the existing extensible attributes:

- For descendants that already have this extensible attribute, the existing extensible attribute will always be set to **Inherit**.
- For descendants that do not have this extensible attribute, the descendants will not inherit this extensible attribute (extensible attribute is set to **Do not Inherit**).

**Result:**

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	DEF	Native
10.0.0.0/16	Network	Region	DEF	Inherited, Source 10.0.0.0/8
10.1.0.0/16	Network	Region		

#### Example 4

Examples in this section show different results when you remove an existing inheritable extensible attribute from an object located at the top or in the middle of the inheritance chain based on the following:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	DEF	Native
10.0.0.0/16	Network	Region	DEF	Inherited from 10.0.0.0/8
10.1.0.0/16	Network	Region	ABC	Overridden

#### Example 4.1: Remove extensible attribute Region with value DEF from 10.0.0.0/8

You select the following option for the existing extensible attribute:

- Remove extensible attributes with inheritance state set to **Inherited**. Extensible attributes with the state set to **Overridden** are not removed.

**Result:**

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region		
10.0.0.0/16	Network	Region		
10.1.0.0/16	Network	Region	ABC	Overridden

**Example 4.2: Remove extensible attribute Region with value DEF from 10.0.0.0/8** You select the following option for the existing extensible attribute:

- Preserve all descendant extensible attributes, whether the state is set to **Inherited** or **Overridden**. **Result:**

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container			
10.0.0.0/16	Network	Region	DEF	Native
10.1.0.0/16	Network	Region	ABC	Native

#### Example 5

Examples in this section show different results when you remove parent object based on the following:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	DEF	Native
10.0.0.0/16	Network Container	Region	DEF	Inherited from 10.0.0.0/8
10.1.0.0/24	Network	Region	DEF	Inherited from 10.0.0.0/8
10.1.0.0/24	Network	Region	ABC	Overriden
10.10.0.0/16	Network Container	Region	GHI	Overriden
10.10.0.0/24	Network	Region	GHI	Inherited from 10.10.0.0/16
10.10.0.0/24	Network	Region	JKL	Overriden

**Example 5.1: Removing object 10.0.0.0/8 from the parent level** You select the following option for the existing extensible attribute:

- Remove extensible attributes with the inheritance state set to **Inherited**. Extensible attributes with the state set to **Overriden** are not removed.

**Result:**

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/16	Network Container			
10.0.0.0/24	Network			
10.0.1.0/24	Network	Region	ABC	Overriden
10.10.0.0/16	Network Container	Region	GHI	Overriden
10.10.0.0/24	Network	Region	GHI	Inherited from 10.10.0.0/16
10.10.1.0/24	Network	Region	JKL	Overriden

**Example 5.2: Removing object 10.0.0.0/8 from the parent level**

You select the following option for the existing extensible attribute on descendants:

- Preserve all extensible attributes on the descendant.

**Result:**

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/16	Network Container	Region	DEF	Native

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/24	Network	Region	DEF	Inherited from 10.0.0.0/16
10.0.1.0/24	Network	Region	ABC	Overridden
10.10.0.0/16	Network Container	Region	GHI	Native
10.10.0.0/24	Network	Region	GHI	Inherited from 10.10.0.0/16
10.10.1.0/24	Network	Region	JKL	Overridden

**Example 5.3: Remove object 10.10.0.0/16 from the parent level**

You select the following option for the existing extensible attribute on descendants:

- Remove extensible attributes with the inheritance state set to **Inherited**. Extensible attributes with the state set to **Overridden** are retained.

**Result:**

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	DEF	Native
10.0.0.0/16	Network Container	Region	DEF	Inherited from 10.0.0.0/8
10.0.0.0/24	Network	Region	DEF	Inherited from 10.0.0.0/8
10.0.1.0/24	Network	Region	ABC	Overridden
10.10.0.0/24	Network	Region		
10.10.1.0/24	Network	Region	JKL	Overridden

**Example 6**

Examples in this section show different results after you add an object in the middle of the inheritance chain based on the following:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Owner	Admin	Native
10.0.0.0/16	Network Container	Owner	Admin	Inherited from 10.0.0.0/8
10.0.0.0/24	Network	Owner	Admin	Inherited from 10.0.0.0/8

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.1.0/24	Network	Owner	Joe	Overridden
10.10.0.0/24	Network	Owner	Admin	Inherited from 10.0.0.0/8
10.10.1.0/24	Network	Owner	Annie	Overridden

**Example 6.1: Adding object 10.10.0.0/16 without extensible attributes** You select the following option for the existing extensible attribute:

- Retain values if the extensible attribute already exists, and inherit the attribute from the parent object if it does not exist.

**Result:**

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Owner	Admin	Native
10.0.0.0/16	Network Container	Owner	Admin	Inherited from 10.0.0.0/8
10.0.0.0/24	Network	Owner	Admin	Inherited from 10.0.0.0/8
10.0.1.0/24	Network	Owner	Joe	Overridden
10.10.0.0/16	Network Container	Owner	Admin	Inherited from 10.0.0.0/8
10.10.0.0/24	Network	Owner	Admin	Inherited from 10.0.0.0/8
10.10.1.0/24	Network	Owner	Annie	Overridden

### Example 7

Examples in this section show different results after you modify inheritable extensible attributes with multiple values based on the following:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container			
10.0.0.0/16	Network Container	Region	MNO	Native
		Region	PQR	Native

**Example 7.1: Adding extensible attribute Region with value GHI to 10.0.0.0/8** You select the following option for the existing extensible attributes:

- The descendants that already have this extensible attribute will inherit the value from the parent object.

**Result:** Multiple values will be replaced with the single inherited value.

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	GHI	Native
10.0.0.0/16	Network Container	Region	GHI	Inherited from 10.0.0.0/8

**Example 7.2: Adding extensible attribute Region with value GHI to 10.0.0.0/8** You select the following option for the existing extensible attributes:

- The descendants that already have this extensible attribute will override the value.

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Region	GHI	Native
10.0.0.0/16	Network	Region	DEF	Overridden
		Region	ABC	Overridden

### Example 8

Examples in this section show different results after you modify existing inheritable extensible attribute of an object, but you do not have required permission to modify some descendants. For information about admin permissions, see [About Administrative Permissions](#).

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State	Permission
10.0.0.0/8	Network Container	Owner	Sam	Native	Write
10.0.0.0/16	Network Container	Owner	Sam	Inherited from 10.0.0.0/8	Read
10.0.0.0/24	Network	Owner	Sam	Inherited from 10.0.0.0/8	Read
10.0.1.0/24	Network	Owner	Bob	Overridden	Write

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State	Permission
10.10.0.0/16	Network Container	Owner	John	Inherited from 10.0.0.0/16	Read
10.10.1.0/24	Network	Owner	Max	Overridden	Read
10.20.0.0/16	Network Container	Owner	Sam	Inherited from 10.0.0.0/8	Write
10.20.0.0/24	Network	Owner			Read
10.20.1.0/24	Network	Owner			Read

**Example 8.1: Removing object 10.0.0.0/8**

You select the following option for the existing inheritable extensible attribute:

- Retain extensible attribute values on descendants that are inherited from this parent object.

**Result:**

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State	Permission
10.0.0.0/16	Network Container	Owner	Sam	Native	Read
10.0.0.0/24	Network	Owner	Sam	Native	Read
10.0.1.0/24	Network	Owner	Bob	Overridden	Write
10.10.0.0/16	Network Container	Owner	John	Overridden	Read
10.10.0.0/24	Network	Owner	John	Inherited from 10.10.0.0/16	Read
10.10.1.0/24	Network	Owner	Max	Overridden	Read
10.20.0.0/16	Network Container	Owner	Sam	Native	Write
10.20.0.0/24	Network				Read
10.20.1.0/24	Network				Read



### Example 8.2: Removing object 10.0.0.0/8

You select the following option for the existing inheritable extensible attribute:

- Remove extensible attribute values from descendants that are inherited from this parent object.

The appliance displays an error message when you remove an extensible attribute that is associated with a descendant for which you do not have required permission.

#### Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State	Permission
10.0.0.0/16	Network Container				
10.0.0.0/24	Network				
10.0.1.0/24	Network	Owner	Bob	Overridden	Write
10.10.0.0/16	Network Container	Owner	John	Overridden	Read
10.10.0.0/24	Network	Owner	John	Inherited from 10.10.0.0/16	Read
10.20.0.0/16	Network	Owner	Max	Overridden	Read
10.20.0.0/24	Network Container				Write
10.20.1.0/24	Network				Read

### Example 9

Examples in this section show different results after you join multiple networks, based on the following:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.1.0.0/16	Network	Owner	John	Native
10.1.0.1	Fixed Address	Owner	John	Inherited from 10.1.0.0/16
10.2.0.0/16	Network	Owner	Sam	Native
10.2.0.1	Fixed Address	Owner	Jane	Overridden

### Example 9.1: Joining networks 10.0.0.0/8 with 10.1.0.0/16

You select the following option for the existing extensible attribute:

- Join networks 10.0.0.0/8 with 10.1.0.0/16.

Result:

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Owner	John	Native
10.1.0.1	Fixed Address	Owner	John	Inherited from 10.0.0.0/8
10.2.0.1	Fixed Address	Owner	Jane	Overridden

### Example 9.2: Joining networks 10.0.0.0/8 with 10.2.0.0/16

You select the following option for the existing extensible attribute:

- Join networks 10.0.0.0/8 with 10.2.0.0/16.

Object	Type	Extensible Attribute	Extensible Attribute Value	Inheritance State
10.0.0.0/8	Network Container	Owner	Sam	Native
10.1.0.1	Fixed Address	Owner	Sam	Inherited from 10.0.0.0/8
10.2.0.1	Fixed Address	Owner	Jane	Overridden

## Managing Security Operations

The Grid provides certain security-related features. The following sections describe the different security-related features that you can set.

### Restricting Remote Console Access

You can restrict admins from accessing the Infoblox CLI from a remote location using an SSH (Secure Shell) v2 client. When you select this only admins who have access rights will be able to access the Infoblox CLI using an SSH (Secure Shell) v2 client. By default, this option is disabled.

### Restricting GUI/API Access

You can specify the IP addresses from which administrators are allowed to access the NIOS appliance. When the NIOS appliance receives a connection request, it tries to match the source IP address in the request with IP addresses in the list. If there is at least one item in the HTTP Access Control list and the source IP address in the request does not match it, the NIOS appliance ignores the request.

---

**Caution:** If you specify an address or network other than the one from which you are currently accessing the appliance, when you save your configuration, you will lose your administrative session and be unable to reconnect. If you have enabled the **Enable GUI/API Access via both MGMT and LAN1/VIP** feature and configured ACLs to control access to the

GUI and API, then the same set of ACLs are applicable on both the interfaces (LAN1 and MGMT port). For information, see [Enabling GUI and API Access on the MGMT and LAN1/VIP Ports](#) and Configuring Security Features below.

---

## Enabling HTTP Redirection

You can enable the NIOS appliance to redirect administrative connection requests using HTTP to the secure HTTPS protocol. When you disable redirection, the NIOS appliance ignores any administrative connection requests not using HTTPS. By default, the NIOS appliance does not redirect HTTP connection requests to HTTPS. When you change this setting, the application restarts and your management session terminates.

## Modifying the Session Timeout Setting

You can set the length of idle time before an administrative session to the Infoblox GUI times out. The default timeout value is 600 seconds (10 minutes).

If a user does not interact with the application for the specified time, the appliance displays a message that a timeout has occurred. Click **OK** to restart the GUI session. Note that the session timeout is not honored when the Infoblox GUI supports auto-refresh. This means that the Infoblox interface page is automatically refreshed even after you have set the session timeout value.



### Warning

If the [Detailed Status Panel](#) is open, the following actions take place:

- Grid Manager auto refreshes at a rate of 30 seconds.
- Widgets that support user-specified auto refresh, refresh at the rate specified in the **Auto Refresh Period** field.

Therefore, even if you set the session timeout to be greater than the auto refresh rate, auto refresh still takes place. The Grid Manager session does not time out because auto refresh takes precedence over the session timeout. For more information about widgets, see [Status Dashboard](#).



### Note

If you change the session timeout value, the new setting takes effect only after you log out and log back in.

## Disabling the LCD Input Buttons

By default, the LCD input function is enabled, which allows you to use the LCD buttons on the front panel of a NIOS appliance to change the IPv4 address settings of the LAN port. You can disable this function if the appliance is in a location where you cannot restrict access exclusively to NIOS appliance administrators and you do not want anyone to be able to make changes through the LCD.

## Configuring Security Features

You can manage only certain features at the member level. To configure security features for the Grid or an individual member:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties** -> **Edit**.  
**Member:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then

click the Edit icon.

To override an inherited property, click **Override** next to it and complete the appropriate fields.

2. On the **Security** tab, complete the following:
  - **Session Timeout(s)**: This field is in the *Grid Properties* editor only. Enter a number between 60 and 31536000 seconds (one minute – one year) in the Session Timeout field. The default session timeout is 600 seconds (10 minutes).
  - **Minimum Password Length**: This field is in the *Grid Properties* editor only. Specify the minimum number of characters allowed for an admin password.
  - **Redirect HTTP to HTTPS**: This field is in the *Grid Properties* editor only. Select this option to have the appliance redirect HTTP connection requests to HTTPS.
  - **Restrict GUI/API Access**: To control access to the GUI and API, select one of the following. You can restrict access using a named ACL or define individual ACEs. For information about named ACLs, see [Configuring Access Control](#).
    - **Allow Any**: Select this to allow any clients to access the GUI and API. This is selected by default.
    - **Named ACL**: Select this and click **Select Named ACL** to select a named ACL that contains only IPv4 and IPv6 addresses and networks. GUI and API access restriction does not support TSIG key based ACEs. When you select this, the appliance allows GUI and API access for all ACEs in the named ACL. You can click **Clear** to remove the selected named ACL.
    - **Set of ACEs**: Select this to configure individual access control entries (ACEs). Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding.
      - **IPv4 Address and IPv6 Address**: Select this to add an IPv4 address or an IPv6 address. Click the **Value** field and enter the IP address. The appliance allows this client to access the GUI and API and restricts others.
      - **IPv4 Network and IPv6 Network**: Select this to add an IPv4 network or IPv6 network. Click the **Value** field and enter the network. The appliance allows this network to access the GUI and API and restricts others.  
After you have added access control entries, you can do the following:
        - Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
        - Reorder the list of ACEs using the up and down arrows next to the table.
        - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
    - **Access Restrictions Apply to Remote Console**: Select this to restrict admins from accessing the Infoblox CLI from a remote location using SSH (Secure Shell) v2 client.
    - **Enable Remote Console Access**: Select this option to enable superuser admins to access the Infoblox CLI from a remote location using SSH (Secure Shell) v2 clients. You can set this at the Grid and member levels.
    - **Enable Support Access**: Select this checkbox to enable an SSH (Secure Shell) daemon that only Infoblox Technical Support can access. You can set this at the Grid and member levels.
    - **Restrict Remote Console and Support Access to the MGMT Port**: This field is in the *Grid Member Properties* editor only. Select this checkbox to restrict SSH (Secure Shell) v2 access to the MGMT port only. This restricts Infoblox Technical Support and remote console connections—both of which use SSH v2—to just the MGMT port. For an HA pair, you can make an SSH v2 connection to the MGMT port on both the active and passive node.  
Clear the checkbox to allow SSH v2 access to both the MGMT and LAN ports.
    - **Permanently Disable Remote Console and Support Access**: This field is in the *Grid Properties* editor only.  
Select this option to permanently disable remote console (Secure Shell v2) access for appliance administration and for Infoblox Technical Support.
    - **Enable LCD Input**: Select this checkbox to allow use of the LCD buttons on the front panel of a NIOS appliance to change the IP address settings of the LAN port. Clear this checkbox to disable this function. You can set this at the Grid and member levels.
    - **Disable Concurrent Login**: Select this checkbox to disallow multiple logins per user for the same NIOS session. That is, if you have already logged on to one NIOS session (for example https://

255.255.255.0) you cannot log on to the same IP address from another browser or from another system.

Note that before you disable multiple logins to a NIOS system, ensure that all its existing sessions (if any) are logged out. If not, the existing sessions continue to remain active even after you disable multiple logins.

- **Enable Account Lockout:** Select the checkbox to enable account lockout for the local user. You can enable password security such that if a local user tries to log in to Grid Manager by using an incorrect password, NIOS appliance locks the user account after the configured number of failed login attempts for a configured time period. Only superusers can enable and configure this feature. This feature is applicable only to local users. This option is disabled by default.
  - **Maximum number of attempts:** Enter the maximum number of invalid login attempts to Grid Manager after which NIOS locks the account. You can specify a value from **1** to **99**. The default value is **5**.
  - **Lockout duration:** Enter the time duration in minutes for which the account must be locked. You can specify a value from **1** to **1440**. The default value is **5 mins**.
  - **Never Unlock:** Select the checkbox to permanently lock a local user account which is already locked. Only a superuser can clear the checkbox to unlock the account. NIOS displays a warning message if you enable this option. This option is not applicable to superuser accounts because you cannot permanently lock a superuser account. This option is disabled by default.

You can also configure account lockout for admin groups. For more information, see [Configuring Account Lockout for Admin Group](#).
- **Disable Inactive Users:** Select this checkbox to disable users who have not logged in to NIOS for a specified period of time. You can specify the time period (in days) in the **Disable account if user has not logged in for <time period> days** field. The range of days is from 2 to 9999. You can also specify a reminder to be sent in the **Remind <days> prior to expiration** field. The range of days is from 1 to 30. The number of days you specify in this field is the time from which users start getting daily email reminders that their account will be disabled. NIOS sends the email reminder only if an email address has been configured for the user.

Select the **Allow user to reactivate account via serial console** and **Allow user to reactivate account via remote console** checkboxes if you want users to activate their account after it has been disabled. To reactivate using the serial console, see [Deploying a Single Independent Appliance](#). To reactivate using the remote console, type `ssh <user name>@<ip address>`.

**Note:** Reactivating the account using the serial console or the remote console is possible only for superusers.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Enabling and Disabling Remote Console and Infoblox Technical Support Access

Infoblox Technical Support might need access to your NIOS appliance to troubleshoot problems. This function enables an SSH (Secure Shell) daemon that only Infoblox Technical Support can access. By default, this option is disabled. This function also makes it possible for a superuser admin to access the Infoblox CLI from a remote location using an SSH (Secure Shell) v2 client. The management system must have an SSH v2 client to use this function. After opening a remote console connection using an SSH client, log in using a superuser name and password. By default, this option is disabled. Note that only superusers can log in to the appliance through a console connection.

You can permanently disable remote console (Secure Shell v2) access for appliance administration and for Infoblox Technical Support to perform remote troubleshooting. Disabling this type of access might be required in a high-security environment.



### Warning

*After permanently disabling remote console and support access, you cannot re-enable them! Not even resetting an appliance to its factory default settings can re-enable them.*

If you have any questions, contact Infoblox Technical Support. To enable or disable remote console and Infoblox technical support access:

1. **Grid:** From the **Grid** tab, select **Grid Manager** tab, expand the Toolbar and click **Grid Properties** -> Edit.  
**Member:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.  
**Independent appliance:** From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties** -> Edit.
2. In the editor, select the **Security** tab -> **Advanced** tab, and then complete the following in the **Remote Console and Infoblox Technical Support Access** section:
  - **Enable Remote Console Access:** Select this checkbox to enable superuser admins to access the Infoblox CLI from a remote location using SSH (Secure Shell) v2 clients. You can set this at the Grid and member levels.
  - **Enable Support Access:** Select this checkbox to enable an SSH (Secure Shell) daemon that only Infoblox Technical Support can access. You can set this at the Grid and member levels.
  - **Support Access Info:** Displays the support access code and the expiration time of the code. Note that the **Enable Support Access** is disabled after the expiration time.
  - **Permanently Disable Remote Console and Support Access:** This field is in the *Grid Properties* editor only. Select this checkbox to permanently disable remote console (Secure Shell v2) access for appliance administration and for Infoblox Technical Support.
3. Save the configuration.

## Configuring Proxy Servers

If your network environment does not allow direct HTTP or HTTPS communication with the Internet through a firewall from a secure location in which the Grid Master or standalone appliance resides, you can configure your appliance to use a proxy server so you can receive automatic updates, such as threat protection rulesets and threat analytics bundles, through this connection. You can also configure a proxy server to perform AWS related communication, such as using a proxy server as the AWS API Proxy, performing vDiscovery on AWS endpoints, and pulling DNS data from Amazon Route 53. For information about AWS deployments, refer to the *Installation Guide for vNIOs for AWS*. For information about vDiscovery, see [Configuring vDiscovery Jobs](#).



### Note

- Configured proxy settings are for the entire Grid. You cannot configure proxy settings for individual members.
- Infoblox supports only basic authentication for vDiscovery jobs performed over a proxy server as not all proxy options used by customers are validated at the Infoblox lab.

Depending on the updates you want to download, you may need to install the respective licenses in your Grid. For example, to download threat protection ruleset updates, the Grid must have the Threat Protection Update license installed. To download threat analytics bundles, you must install the Threat Analytics license. When you configure your appliance to obtain periodic ruleset updates, all updates go through the MGMT port on the Grid Master by default. You can, however, delegate this function to a Grid member using a different interface such as LAN1 or LAN2. For information about how to delegate updates to a Grid member and configure the interface, see [Configuring Members and Interfaces for Automatic Updates](#) below.

To configure proxy settings for the Grid, complete the following steps:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Edit** -> **Grid Properties** from the Toolbar.
2. In the *Grid Properties* editor, select the **Proxy Settings** tab -> **Basic** tab, and complete the following:
  - **Use Proxy Server:** When you select this checkbox, the appliance uses the connection that has been established with the proxy server to establish connection with endpoints or download automatic updates, such as threat protection rulesets and threat analytics bundles. The reporting member sends API requests to the proxy server for threat details. For more information, see [Threat Protection Reports](#). Similarly, the Grid Master sends API requests to the proxy server for all threat context details. For more information,

see [Viewing the RPZ Threat Details](#). This setting applies to the entire Grid. When you clear this checkbox, the appliance does not use the proxy server; however, the configuration will not be affected.

- **Name or IP Address** and **Port**: Enter the name or IP address and port number of the proxy server you plan to use for this connection.
- **HTTPS Proxy Content Inspection**: From the drop-down list, select one of the following methods the proxy server uses to inspect packet content. Note that this section does not apply to AWS deployments.
  - **None**: Select this to use HTTP for the connection. This method does not allow certificate authentication for the proxy server.
  - **Allow Deep Packet Inspection**: This option is not supported for AWS deployments. To eliminate man-in-the-middle attacks, select this to allow deep pack inspection and information extraction for non-compliant protocol, intrusions, or other criteria that determine whether the packets should be routed to an alternate destination. When you select this, you must click **Proxy Server Certificate** and navigate to the proxy server certificate to upload it to the Grid, or you must ensure that a trusted chain has been established before the proxy server can perform deep packet inspection. When you have uploaded a certificate, the appliance displays **Loaded**.
    - **Enable Strict Host Name Checking**: This option is enabled only when you select **Allow Deep Packet Inspection**. As part of the SSL handshake process, the appliance verifies that the CN (Common Name) of the public certificate of the proxy server exactly matches the host name of the proxy server.

#### Credentials for Proxy Server (if configured at proxy server)

- **Use user name and password to connect to proxy server if configured**: If you have configured user credentials on the proxy server, enter the **Username** and **Password** here. This is optional.

#### Configuring Members and Interfaces for Automatic Updates

If you want to download specific rulesets or updates from external servers, you can configure members and corresponding interfaces to receive updates automatically. Note that these members must have access to external servers from where the updates are delegated automatically. For a member to access an external server, you must specify the interface that you configure here when you set up an external server.

To delegate automatic updates to a Grid member or change the interface for downloads, complete the following steps:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Edit -> Grid Properties** from the Toolbar.
2. In the *Grid Properties* editor, select the **Proxy Settings** tab -> **Basic** tab, and complete the following:
  - In the Members table, click the Add icon and select **Add Member**. In the *Member Selector*, select the member to which you want to delegate the automatic update function. The appliance adds the Grid member to the table. You can add up to five members for automatic updates. When you add multiple members, you can place the selected members in the order you prefer using the arrows next to the table. If for any reasons the first member is unable to receive updates, the appliance goes down the list in the order the members are placed until the next reachable member receives the updates.
  - Click the **Interface** column, and from the drop-down list select the interface (**Any**, **LAN2**, **MGMT** or **VIP**) you want the corresponding member to use for automatic updates. Ensure that the selected interface is configured properly on the member. The default is **Any**, which maps to the LAN1 interface for all appliance models, except for the PT models. For all PT appliances, **Any** maps to the MGMT interface. You can select **MGMT** or **Any** for threat protection members. Note that VLANs are not supported. You can click the **Test connectivity to all members** icon to ensure that the connectivity between the ruleset download web site and all Grid members is properly established using the selected interface. The appliance displays a message indicating the connectivity status in the feedback panel.

The table displays the following:

- **Name**: Displays the member name. For the Grid Master, this field displays **Grid Master**.
- **Interface**: Displays the selected interface that is used for automatic updates.



#### Note

The appliance generates an SNMP trap if any of the configured Grid members failed to receive updates.



## Configuring Ethernet Ports

Depending on your deployment and configuration choices, the Ethernet ports on the NIOS appliance perform different functions. The Ethernet ports that handle traffic on the NIOS appliance are as follows:

- LAN1 port – A 10/100/1000-Mbps gigabit Ethernet port that connects the appliance to the network. This is the default port for single independent appliances, single Grid members, and passive nodes in HA pairs. You must use the LAN1 port to set up the appliance initially. It handles traffic for all management services if you do not enable the MGMT and LAN2 ports. The passive node in an HA pair uses this port to synchronize the database with the active node.
  - LAN2 port – A 10/100/1000-Mbps gigabit Ethernet port that connects the appliance to the network. The LAN2 port is not enabled by default. You can enable the LAN2 port and define its use through the GUI after the initial setup. By default, the appliance uses the LAN1 port (and HA port when deployed in an HA pair). To enable and configure the LAN2 port, you must have read/write permission to the Grid member on which you want to enable the port. The LAN2 port is available on the TE-815, TE-825, TE-1415, TE-1425, TE-2215, TE-2225, TE-4015 and TE-4025 appliances. For information about how to use the LAN2 port, see [Using the LAN2 Port](#).
  - HA port – A 10/100/1000-Mbps gigabit Ethernet port through which the active node in an HA (high availability) pair connects to the network using a VIP (virtual IP) address. HA pair nodes also use their HA ports for VRRP (Virtual Router Redundancy Protocol) advertisements.
- MGMT port – A 10/100/1000-Mbps gigabit Ethernet port that you can use for appliance management or DNS service. You can enable the MGMT port and define its use through the GUI after the initial setup. If the MGMT port is enabled, the NIOS appliance uses it for management services (see the Sources and Destinations for Services table below for specific types).

You can do the following on some of the Ethernet ports, depending on your network requirements and configurations:

- Assign VLANs (Virtual LANs) to the LAN1 and LAN2 ports so that NIOS can provide DNS service to different subnetworks on the same interface.
- Implement DiffServ (Differentiated Services) on the appliance by configuring the DSCP (Differentiated Services Code Point) value.

### Enabling GUI and API Access on the MGMT and LAN1/VIP Ports

You can access the Infoblox GUI and API through the MGMT and LAN1 or VIP interfaces simultaneously. To do so, you must first configure the MGMT port on the appliance, and then enable the **Enable GUI/API Access via both MGMT and LAN1/VIP** feature. For information about the MGMT port, see [Using the MGMT Port](#). When you enable this feature, you can use the MGMT and LAN1 ports for standalone appliances and MGMT and VIP ports for an HA pair. This feature is disabled for all new installations and upgrades.



#### Note

When the Threat Protection service is running on the Advanced Standalone Appliance, then the GUI and API access is allowed only on the MGMT port.

To enable GUI and API access on the MGMT and LAN1/VIP ports:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and select **Grid Properties** -> **Edit**.
3. In the *Grid Properties* editor, select the **General** tab -> click the **Advanced** tab (or click Toggle Advanced Mode) and complete the following:
  - **Enable GUI/API Access via both MGMT and LAN1/VIP**: Select this checkbox to allow access to the Infoblox GUI and API using both the MGMT and LAN1 ports for standalone appliances and allow both the MGMT and VIP ports for an HA pair. This feature is valid only if you have enabled the MGMT port. For information about enabling the MGMT port, see .
4. Click Save to save the changes.



## About Virtual LANs

You can assign VLANs (Virtual Local Area Networks) to the LAN1, LAN2, and VIP (for HA pairs) interfaces so the appliance can provide DNS service to different subnetworks on the same interface. You can also configure VLANs interfaces on supported Network Insight appliances and use them exclusively for discovery purposes. VLANs are independent logical networks that are mutually isolated on the interface so that IP packets can pass between them through one or more switches or routers. You can assign VLANs to provide segmentation services to address issues such as scalability, security, and network management. For example, you can partition your network into segments such as DHCP address allocation, DNS service, guest network, and DMZ (demilitarized zone) to achieve a higher level of security and to increase performance by limiting broadcast domains. You can also add quality of service schemes to optimize your network traffic on the VLAN trunk links by configuring the DSCP (Differentiated Services Code Point) value for the corresponding physical and virtual interfaces.



### Note

When you configure VLANs on the following Network Insight appliances: ND-1405, ND-2205, ND-4000, ND-V1405, and ND-V2205, the VLAN interfaces are used exclusively for discovery. You cannot bind other services on these VLAN interfaces of the supported Network Insight appliances. For more information about Network Insight, see [About Network Insight](#).

## VLAN Tagging

When your VLANs span across multiple networks, VLAN tagging is required. This enables the NIOS appliance to connect to different networks using the same port. VLAN tagging involves adding a VLAN tag or ID to the header of an IP packet so the appliance can identify the VLAN to which the packet belongs. In addition, switches use the VLAN tag to determine the port to which it should send a broadcast packet. The appliance uses the IEEE 802.1Q networking standard to support VLANs and VLAN tagging. On the appliance, you can configure VLANs as tagged networks by adding VLAN tags to them. You can create up to 10 IPv4 and IPv6 addresses per interface and configure a VLAN ID from one to 4094. You can also configure an address, gateway, and a netmask for VLAN. Any IPv4 or IPv6 address with a VLAN ID is considered as a tagged network. For HA pairs, the appliance supports only one VLAN interface for VRRP over an IPv4 or IPv6. It supports one untagged IPv4 and IPv6 address for each interface and considers this as the primary IP address for the network. For an HA pair, if you have multiple VLANs assigned to a VIP interface, then a network failure in any one of the VLAN interface does not trigger a failover of the active member.

Untagged networks are those without VLAN tags assigned to them. When you set up a VLAN as either a tagged or untagged network, ensure that you properly configure the corresponding switch for the VLAN to function properly.



### Note

A tagged VLAN interface receives only those packets that belongs to the tagged network, but an untagged VLAN interface receives all the packets belonging to the tagged and untagged networks of the interface.

VLANs and VLAN tagging are supported on both IPv4 and IPv6 transports. This feature is currently supported on the following Infoblox appliances: Trinzic 1405, 1415, 1425, 2205, 2215, 2225, 4005, Infoblox-4030-10GE, PT-1405, PT-2205, CP-VM-800, CP-VM-1400, and CP-VM-2200. It is also supported on all the Trinzic virtual appliances. VLAN tagging is not supported on TE-100, TE-805, ND-805, TR-805, TE-815, and TE-825. For information about these appliances, refer to the respective installation guides on the Infoblox Support web site at <https://www.infoblox.com/support>.

Currently, only the DNS service can listen on specific VLAN interfaces. The DHCP service listens only on the primary VLAN interface (tagged or untagged). You can also specify VLANs as the source port for sending DNS queries and notify messages. For information about how to configure these, see [Specifying Port Settings for DNS](#).

Additional VLAN support is available exclusively for discovery on the following Network Insight appliances: ND-1405,

ND-2205, ND-4000, ND-V1405, and ND-V2205. Binding other services on the VLAN interfaces of the Network Insight appliances is not supported.



#### Note

When you join an appliance that supports VLANs to a Grid that does not support VLANs or revert the appliance to a NIOS version that does not support VLANs, the appliance will become unreachable after joining the Grid or being reverted. You must remove VLAN tagging from the corresponding switch in order to reach the downgraded appliance.

Consider the following guidelines when tagging VLANs on the LAN1 and LAN2 ports:

- You can assign VLAN addresses to an interface and add VLAN tags to them. However, you must designate one of the tagged VLANs as a primary address.
- If the primary IPv4 address is tagged with a VLAN ID, all other addresses on the same interface must be tagged with a VLAN ID as well.
- You can use the same VLAN ID to tag only one IPv4 and one IPv6 address on the same interface. You cannot use the same VLAN ID to tag multiple IPv4 and IPv6 addresses on the same interface.
- You can assign one untagged IPv4 and one untagged IPv6 address to the same interface. These addresses are designated as the primary address for the interface.
- For IPv6, you must have a primary IPv6 address (either tagged or untagged) before you can add other tagged IPv6 addresses on the same interface.
- If you have multiple VLANs assigned to the LAN1 interface and the primary VLAN is untagged, DHCP listens on all VLAN interfaces and thus DHCP lease requests will succeed for the additional VLANs assigned to the LAN1 interface, but the request will actually be handled by the primary untagged VLAN interface.
- You can set up the system to define only tagged networks:
  - When the VLAN tag is not set, the appliance considers the network as an untagged network.
  - You can specify a single untagged IPv4 and IPv6 network per interface.
  - The primary network can be tagged or untagged, but you must tag the additional VLANs.
- VMware ESXi does not support tagged and untagged subnets on the same interface simultaneously. As a result, VMware vSwitch cannot perform tag translation on a tagged VLAN ID if there is another untagged VLAN on the same interface.

## Configuring VLANs

When you first set up a NIOS appliance, you can assign VLANs through the Grid Setup Wizard. After the initial setup, you can assign VLANs to the LAN1 or LAN2 ports in the Required Ports and Addresses table, as described in Modifying Ethernet Port Settings.

On a Grid member, you can assign up to 10 VLANs for each protocol (IPv4 or IPv6) on the LAN1 and LAN2 ports. You can assign up to 10 IPv4 VLAN addresses and 10 IPv6 VLAN addresses for each interface. You can configure only IPv4 VLAN addresses for an IPv4 Grid member and only IPv6 VLAN addresses for an IPv6 Grid member, but for a dual mode Grid member you can configure both IPv4 and IPv6 VLAN addresses.

To assign additional VLANs to the LAN1 or LAN2 port, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. Select the **Network** -> **Basic** tab in the *Grid Member Properties* editor.
3. In the **Additional Ports and Addresses** table, click the Add icon and select either **MGMT (IPv4)**, **MGMT (IPv6)**, **LAN2 (IPv4)**, **LAN2 (IPv6)**, **Additional Address (loopback) (IPv4)**, **Additional Address (loopback) (IPv6)**, **LAN1 (VLAN)(IPv4)**, **LAN1 (VLAN)(IPv6)**, **LAN2 (VLAN)(IPv4)** or **LAN2 (VLAN)(IPv6)** from the drop-down list. You can add up to 10 IPv4 and 10 IPv6 VLANs for each interface.  
You can configure only IPv4 VLAN addresses for an IPv4 Grid member and only IPv6 VLAN addresses for an IPv6 Grid member, but for a dual mode Grid member you can configure both IPv4 and IPv6 VLAN addresses.
  - For vNIOS appliances, some of the options in the drop-down list may vary depending on your vNIOS configuration. For example, if you are using a single network interface instance of vNIOS for GCP, you will see choices specific to the LAN1 interface and Additional Address only. For more information, see the vNIOS documentation specific to your product at [Appliances](#). Note that you can:

- You can configure only IPv4 VLAN addresses for an IPv4 Grid member and only IPv6 VLAN addresses for an IPv6 Grid member, but for a dual mode Grid member you can configure both IPv4 and IPv6 VLAN addresses.
  - For vNIOS appliances, some of the options in the drop-down list may vary depending on your vNIOS configuration. For example, if you are using a single network interface instance of vNIOS for GCP, you will see choices specific to the LAN1 interface and Additional Address only. For more information, see the vNIOS documentation specific to your product at [Appliances](#).
  - **MGMT (IPv4)**: Select this to configure IPv4 address for MGMT port.
  - **MGMT (IPv6)**: Select this to configure IPv6 address for MGMT port.
  - **LAN2 (IPv4)**: Select this to configure IPv4 address for the LAN2 port for DHCP or DNS. This is not applicable to Trinzic 100 appliance.
  - **LAN2 (IPv6)**: Select this to configure IPv6 address for the LAN2 port for DHCP or DNS. This is not applicable to Trinzic 100 appliance.
  - **Additional Address (loopback) (IPv4)**: Select this to add a non-anycast IPv4 address to the loopback interface. Note that you can configure this for IPv4 and dual mode Grid member.
  - **Additional Address (loopback) (IPv6)**: Select this to add a non-anycast IPv6 address to the loopback interface. Note that you can configure this for IPv6 and dual mode Grid member.
  - **LAN1 (VLAN) (IPv4)**: Select this to add a VLAN to the LAN1 interface. You can add up to 10 IPv4 VLAN addresses. Note that you can configure this for IPv4 and dual mode Grid member. This is supported on the following Infoblox appliances: Trinzic 1405, 1415, 1425, 2205, 2215, 2225, 4005, Infoblox-4030-10GE, PT-1405, PT-2205, CP-VM-800, CP-VM-1400, and CP-VM-2200. It is also supported on all the Trinzic virtual appliances. VLAN tagging is not supported on TE-100, TE-805, ND-805, TR-805, TE-815, and TE-825.
  - **LAN1 (VLAN) (IPv6)**: Select this to add a VLAN to the LAN1 interface. You can add up to 10 IPv4 and 10 IPv6 VLAN addresses. Note that you can configure this for IPv6 and dual mode Grid member. This is supported on the following Infoblox appliances: Trinzic 1405, 1415, 1425, 2205, 2215, 2225, 4005, Infoblox-4030-10GE, PT-1405, PT-2205, CP-VM-800, CP-VM-1400, and CP-VM-2200. It is also supported on all the Trinzic virtual appliances. VLAN tagging is not supported on TE-100, TE-805, ND-805, TR-805, TE-815, and TE-825.
  - **LAN2 (VLAN) (IPv4)**: Select this to add a VLAN to the LAN2 interface. You can add up to 10 IPv4 VLAN addresses. Note that you can configure this for IPv4 and dual mode Grid member. This is supported on the following Infoblox appliances: Trinzic 1405, 1415, 1425, 2205, 2215, 2225, 4005, Infoblox-4030-10GE, PT-1405, PT-2205, CP-VM-800, CP-VM-1400, and CP-VM-2200. It is also supported on all the Trinzic virtual appliances. VLAN tagging is not supported on TE-100, TE-805, ND-805, TR-805, TE-815, and TE-825.
  - **LAN2 (VLAN) (IPv6)**: Select this to add a VLAN to the LAN2 interface. You can add up to 10 IPv6 VLAN addresses. Note that you can configure this for IPv6 and dual mode Grid member. This is supported on the following Infoblox appliances: Trinzic 1405, 1415, 1425, 2205, 2215, 2225, 4005, Infoblox-4030-10GE, PT-1405, PT-2205, CP-VM-800, CP-VM-1400, and CP-VM-2200. It is also supported on all the Trinzic virtual appliances. VLAN tagging is not supported on TE-100, TE-805, ND-805, TR-805, TE-815, and TE-825.
4. Enter the following:
- **Interface**: Displays the name of the VLAN interface. This can be **LAN1 (VLAN)(IPv4)**, **LAN1 (VLAN)(IPv6)**, **LAN2 (VLAN)(IPv4)**, or **LAN2 (VLAN)(IPv6)** depending on your selection. You cannot modify this.
  - **Address**: Type the IP address for the VLAN port.
  - **Subnet Mask (IPv4) or Prefix Length (IPv6)**: For IPv4 address, specify an appropriate subnet mask and for IPv6 address, specify the prefix length. The prefix length ranges from 2 to 127, with common-sense values ranging from /48 to /127 due to the larger number of bits in the IPv6 address.
  - **Gateway**: Type the IPv4 or IPv6 default gateway address for the VLAN port depending on the type of interface. For IPv6 interface, you can also type **Automatic** to enable the appliance to acquire the IPv6 address of the default gateway and the link MTU from router advertisements.  
You can now define a link-local address as the default IPv6 gateway and isolate the LAN segment so the local router can provide global addressing and access to the network and Internet. This is supported for both LAN1 and LAN2 interfaces as well as LAN1 and LAN2 in the failover mode.
  - **VLAN Tag**: Enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
  - **Port Settings**: For IPv4 only. From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to

- negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
- **DSCP Value:** Displays the Grid DSCP value, if configured. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Implementing Quality of Service Using DSCP

You can implement DiffServ (Differentiated Services) on the appliance by configuring the DSCP (Differentiated Services Code Point) value. DiffServ is a scalable and class-based mechanism that provides relative priorities to the type of services on your network. It can provide low latency for critical network traffic while providing simple best-effort service for non-critical services. The Infoblox DSCP implementation fully conforms to RFC 2475. For more information about DiffServ, refer to RFC 2475, *An Architecture for Differentiated Services*.

In IPv4 and IPv6 headers, DiffServ uses the DS (Differentiated Services) field for packet classification purposes. The DS field defines the layout of the ToS (Type of Services) octet in IPv4 and the Traffic Class octet in IPv6. The first six bits of the DS field are used as the DSCP value, which determines the PHBs (per-hop behaviors) on DiffServ compliant nodes and enables priorities of services to be assigned to network traffic. For more information about the DS field, refer to RFC 2474, *Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers*.

When you configure the DSCP value for DiffServ, the appliance sets priorities for all outgoing IP traffic. It implements QoS (quality of service) rules so you can effectively classify and manage your critical network traffic. To ensure that core network services, such as DNS services, continue to operate in the event of network traffic congestion, you can set the DSCP value for the entire Grid and override it at the member level. Note that on an appliance, all outgoing IP traffic on all interfaces uses the same DSCP value.

DSCP is supported on both IPv4 and IPv6 transports and the DSCP value for both IPv4 and IPv6 transports must be the same. This feature is currently supported on the following Infoblox appliances: Trinzic 2215, 2225, Infoblox-4030-10GE, PT-1405, PT-2205, TE-1415, TE-1425, and TE-4015. For information about these appliances, refer to the respective installation guides on the Infoblox Support web site at <https://www.infoblox.com/support>.



### Note

- DSCP is not supported when packets are processed by DNS cache acceleration over software-based DNS cache acceleration appliances: IB-22x5, IB-v22x5, IB-40x5, IB-v40x5.
- You can set the DSCP value of the primary LAN using the [set network](#) CLI command. DSCP values for all other interfaces and VLANs must be set through Grid Manager.

## Configuring the DSCP Value

The DSCP value is set to zero (lowest priority) by default. You can change this value for the Grid and override the value at the member level. When you configure the DSCP value at the Grid or member level, all outgoing IP traffic on all interfaces uses the same value. Valid DSCP values are from 0 to 63. You can also set the DSCP value using the Infoblox CLI. For more information, refer to the *Infoblox CLI Guide*.

To configure the DSCP value for the Grid:

1. From the **Grid** tab -> **Grid Manager** tab, click **Grid Properties** -> **Edit** from the toolbar.
2. In the **General** -> **Advanced** tab of the *Grid Properties* editor, complete the following:
  - **DSCP Value:** Enter a value from 0 to 63. The default is 0 and it represents the lowest priority.
3. Save the configuration.

To override the DSCP value for a member:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. In the **Network** tab -> **Basic** tab of the *Grid Member Properties* editor, complete the following:

- **DSCP Value:** Click **Override**, and then enter a value from 0 to 63. The default is 0 and it represents the lowest priority.

3. Save the configuration.

You can override the Grid and member DSCP value at the interface level. For more information, see the following:

- For the LAN1 port, see [Modifying Ethernet Port Settings](#) below.
- For the LAN2 port, see [Using the Lan2 Port](#).
- For the MGMT port, see [Using The MGM Port](#).

## Ethernet Port Usage

This section provides tables that detail the port usage and source and destination ports for different services, depending on your Grid configuration. The below table displays the type of traffic per port for both Grid and independent deployments. For a more detailed list of the different types of traffic, see the Sources and Destinations for Services table below. If necessary, to modify the port numbers, see

*Appliance Roles and Configuration, Communication Types, and Port Usage*

Appliance Role	HA Pair	HA Status	MGMT Port	Database Synchronization	Core Network Services	Management Services	GUI Access
HA Grid Master	Yes	Active	Disabled	VIP on HA	VIP on HA	LAN1	VIP on HA
HA Grid Master	Yes	Passive	Disabled	LAN1	–	LAN1	–
Single Grid Master	No	–	Disabled	LAN1	LAN1	LAN1	LAN1
HA Grid Member	Yes	Active	Disabled	LAN1	VIP on HA	LAN1	–
HA Grid Member	Yes	Passive	Disabled	LAN1	–	LAN1	–
Single Grid Member	No	–	Disabled	LAN1	LAN1	LAN1	–
Independent HA Pair	Yes	Active	Disabled	VIP on HA	VIP on HA	LAN1	VIP on HA
Independent HA Pair	Yes	Passive	Disabled	LAN1	–	LAN1	–
Single Independent	No	–	Disabled	–	LAN1	LAN1	LAN1
HA Grid Master	Yes	Active	Enabled	VIP on HA	VIP on HA	MGMT	MGMT
HA Grid Master	Yes	Passive	Enabled	LAN1	–	MGMT	–
Single Grid Master	No	–	Enabled	LAN1	LAN1 or MGMT	MGMT	MGMT and LAN1/ VIP
HA Grid Member	Yes	Active	Enabled	LAN1 or MGMT	VIP on HA	MGMT	–

Appliance Role	HA Pair	HA Status	MGMT Port	Database Synchronization	Core Network Services	Management Services	GUI Access
HA Grid Member	Yes	Passive	Enabled	LAN1 or MGMT	–	MGMT	–
Single Grid Member	No	–	Enabled	LAN1 or MGMT	LAN1 or MGMT	MGMT	–
Independent HA Pair	Yes	Active	Enabled	VIP on HA	VIP on HA	MGMT	MGMT
Independent HA Pair	Yes	Passive	Enabled	LAN1	–	MGMT	–
Single Independent	No	–	Enabled	–	LAN1 or MGMT	MGMT	MGMT
Reporting Member	No	–	Enabled	LAN1 or MGMT	LAN1 or MGMT	MGMT	MGMT

*Appliance Roles and Configuration, Communication Types, and Port Usage for Appliances with LAN2 Ports*

Appliance Role	HA Status	MGMT Port	LAN2 Port	Database Synchronization	Core Network Services	Management Services	GUI Access
HA Grid Master	Active	Disabled	Enabled	VIP on HA	VIP on HA	LAN1 or LAN2	VIP on HA
HA Grid Master	Passive	Disabled	Enabled	LAN1	–	LAN1 or LAN2	–
Single Grid Master	–	Disabled	Enabled	LAN1	LAN1 and/or LAN2	LAN1 or LAN2	LAN1
HA Grid Member	Active	Disabled	Enabled	LAN1	VIP on HA	LAN1 or LAN2	–
HA Grid Member	Passive	Disabled	Enabled	LAN1	–	LAN1 or LAN2	–
Single Grid Member	–	Disabled	Enabled	LAN1	LAN1 and/or LAN2	LAN1 or LAN2	–
Independent HA Pair	Active	Disabled	Enabled	VIP on HA	VIP on HA	LAN1 or LAN2	VIP on HA
Independent HA Pair	Passive	Disabled	Enabled	LAN1	–	LAN1 or LAN2	–
Single Independent	–	Disabled	Enabled	–	LAN1 and/or LAN2	LAN1 or LAN2	LAN1
HA Grid Master	Active	Enabled	Enabled	VIP on HA	VIP on HA	MGMT	MGMT

Appliance Role	HA Status	MGMT Port	LAN2 Port	Database Synchronization	Core Network Services	Management Services	GUI Access
HA Grid Master	Passive	Enabled	Enabled	LAN1	–	MGMT	–
Single Grid Master	–	Enabled	Enabled	LAN1	LAN1, LAN2 and/or MGMT	MGMT	MGMT
HA Grid Member	Active	Enabled	Enabled	LAN1 or MGMT	VIP on HA	MGMT	–
HA Grid Member	Passive	Enabled	Enabled	LAN1 or MGMT	–	MGMT	–
Single Grid Member	–	Enabled	Enabled	LAN1 or MGMT	LAN1, LAN2 and/or MGMT	MGMT	–
Independent HA Pair	Active	Enabled	Enabled	VIP on HA	VIP on HA	MGMT	MGMT
Independent HA Pair	Passive	Enabled	Enabled	LAN1	–	MGMT	–
Single Independent	–	Enabled	Enabled	–	LAN1, LAN2 and/or MGMT	MGMT	MGMT
Reporting Member	–	Enabled	Enabled	LAN1 or MGMT	LAN1, LAN2, and/or MGMT	MGMT	MGMT

To see the service port numbers and the source and destination locations for traffic that can go to and from a NIOS appliance, see the Sources and Destinations for Services table below. This information is particularly useful for firewall administrators so that they can set policies to allow traffic to pass through the firewall as required.



**Note**

The colors in the tables represent a particular type of traffic and correlate with each other.

*Sources and Destinations for Services*

Service	SRC IP	DST IP	Proto	SRC Port	DST Port	Notes
Key Exchange (Member Connection)	LAN1 or MGMT on all Grid members (including Grid Master and Grid Master Candidate)  VIP on HA Grid Master Candidate, or LAN1 on single Grid Master Candidate	VIP on HA Grid Master, or LAN1 on single Grid Master  VIP on HA Grid Master Candidate, or LAN1 on single Grid Master Candidate	17 UDP	2114	2114	Initial key exchange for establishing VPN tunnels Required for Grid
Key Exchange (Grid Master Candidate Promotion)	VIP on HA Grid Master, or LAN1 on single Grid Master  VIP on HA Grid Master Candidate or LAN1 on Single Grid Master Candidate	LAN1 or MGMT on all Grid members (including Grid Master and Grid Master Candidate)	17 UDP	2114	2114	
Accounting	LAN1 or MGMT on Grid member	VIP on HA Grid Master, or LAN1 on single Grid Master VIP on HA Grid Master Candidate, or LAN1 on single Grid Master Candidate	17 UDP	1194 or 5002, or 1024 -> 63999	1194 or 5002, or 1024 -> 63999	Default VPN port 1194 for Grids with new DNSone 3.2 installations and 5002 for Grids upgraded to DNSone 3.2; the port number is configurable  Required for Grid
Network Insight VPN	LAN1 or LAN2 on Probes	LAN1 or LAN2 on Consolidator	UDP	1194	1194	All default VPN tunnels for Network Insight
Discovery	LAN1 or LAN2 on Probes		UDP		161	SNMP
Discovery	LAN1 or LAN2 on Probes		UDP		260	SNMP - Needed for full discovery of some older Check Point models
Discovery	LAN1 or LAN2 on Probes		ICMP		n/a	Ping Sweep



Service	SRC IP	DST IP	Proto	SRC Port	DST Port	Notes
Discovery	LAN1 or LAN2 on Probes		UDP, TCP		53	DNS
Discovery	LAN1 or LAN2 on Probes		ICMP			Path Collection, for IPv4 addresses
Discovery	LAN1 or LAN2 on Probes		UDP		33434+1 per probe packet	Path Collection. Standard traceroute, for IPv6 addresses
Discovery	LAN1 or LAN2 on Probes		ICMP, UDP, TCP			Port scan - all configured by us
Discovery	LAN1 or LAN2 on Probes		UDP		137	NetBIOS
Discovery	LAN1 or LAN2 on Probes		UDP		40125	NMAP, UDP Ping, and credential checking
Discovery	LAN1 or LAN2 on Probes		TCP		23	Telnet can be used based on Network Insight configuration for Network Discovery.
Discovery	LAN1 or LAN2 on Probes		TCP		22	SSH can be used based on Network Insight configuration for Network Discovery.
DHCP	Client	LAN1, LAN2, VIP, or broadcast on NIOS appliance	17 UDP	68	67	Required for IPv4 DHCP service
DHCP	LAN1, LAN2 or VIP on NIOS appliance	Client	17 UDP	67	68	Required for IPv4 DHCP service
DHCP	Client	LAN1, LAN2, VIP, or broadcast on NIOS appliance	17 UDP	546	547	Required for IPv6 DHCP service
DHCP	LAN1, LAN2 or VIP on NIOS appliance	Client	17 UDP	547	546	Required for IPv6 DHCP service

Service	SRC IP	DST IP	Proto	SRC Port	DST Port	Notes
DHCP Failover	LAN1, LAN2 or VIP on Infoblox DHCP failover peer	LAN1, LAN2 or VIP on Infoblox DHCP failover peer	6 TCP	1024 → 65535	519 or 647	Required for DHCP failover
DHCP Failover	VIP on HA Grid Master or LAN1 or LAN2 on single master	LAN1, LAN2 or VIP on Grid member in a DHCP failover pair	6 TCP	1024 -> 65535	647 or 7911	Required for DHCP failover Port 7911 is used by an API for limited control over ISC DHCP server operations.
DDNS Updates	LAN1, LAN2, or VIP	LAN1, LAN2, or VIP	17 UDP	1024 → 65535	53	Required for DHCP to send DNS dynamic updates
DNS Transfers	LAN1, LAN2, VIP, or MGMT, or client	LAN1, LAN2, VIP, or MGMT	6 TCP	53, or 1024 -> 65535	53	For DNS zone transfers, large client queries, and for Grid members to communicate with external name servers Required for DNS
DNS Queries	Client	LAN1, LAN2, VIP, or broadcast on NIOS appliance	17 UDP	53, or 1024 → 65535	53	For DNS queries Required for DNS
DNS Queries	Client	LAN1, LAN2, VIP, or broadcast on NIOS appliance	6 TCP	53, or 1024 → 65535	53	For DNS queries Required for DNS
NTP	NTP client	LAN1, LAN2, VIP, or MGMT	17 UDP	1024 -> 65535	123	Required if the NIOS appliance is an NTP server

Service	SRC IP	DST IP	Proto	SRC Port	DST Port	Notes
NTP	NTP client	LAN1, LAN2, VIP, or MGMT	17 UDP	1024 -> 65535	123	Required if the NIOS appliance is an NTP server. On an HA member, the NTP service runs on the active node. If there is an HA failover, the NTP service is automatically launched after the passive node becomes active and the NTP traffic uses the LAN2, VIP, or MGMT port on one of the nodes from an HA pair, instead of the LAN1 port. During another HA failover, the currently passive node becomes active again and the NTP traffic uses the LAN1 port, and the NTP is back in synchronization.
RADIUS Authentication	NAS (network access server)	LAN1 or VIP	17 UDP	1024 – 65535	1812	For proxying RADIUS Authentication-Requests. The default destination port number is 1812, and can be changed to 1024 – 63997. When configuring an HA pair, ensure that you provision both LAN IP addresses on the RADIUS server.
RADIUS Accounting	NAS (network access server)	LAN1 or VIP	17 UDP	1024 – 65535	1813	For proxying RADIUS Accounting-Requests. The default destination port number is 1813, and can be changed to 1024 – 63998.
RADIUS Proxy	LAN1 or VIP	RADIUS home server	17 UDP	1814	1024 -> 63997 (auth), or 1024 -> 63998 (acct)	Required to proxy requests from RADIUS clients to servers. The default source port number is 1814, and although it is not configurable, it is always two greater than the port number for RADIUS authentication.

Service	SRC IP	DST IP	Proto	SRC Port	DST Port	Notes
ICMP Dst Port Unreachable	VIP, LAN1, LAN2, or MGMT, or UNIX-based client	LAN1, LAN2, or UNIX-based client	1 ICMP Type 3	–	–	Required to respond to the UNIX-based traceroute tool to determine if a destination has been reached
ICMP Echo Reply	VIP, LAN1, LAN2, or MGMT, or client	VIP, LAN1, LAN2, or MGMT, or client	1 ICMP Type 0	–	–	Required for response from ICMP echo request (ping)
ICMP Echo Request	VIP, LAN1, LAN2, or MGMT, or client	VIP, LAN1, LAN2, or MGMT, or client	1 ICMP Type 8	–	–	Required to send pings and respond to the Windows-based traceroute tool
ICMP TTL Exceeded	Gateway device (router or firewall)	Windows client	1 ICMP Type 11	–	–	Gateway sends an ICMP TTL exceeded message to a Windows client, which then records router hops along a data path
NTP	LAN1 on active node of Grid Master or LAN1 of independent appliance	NTP server	17 UDP	1024 -> 65535	123	Required to synchronize Grid, TSIG authentication, and DHCP failover  Optional for synchronizing logs among multiple appliances
SMTP	LAN1, LAN2, or VIP	Mail server	6 TCP	1024 → 65535	25	Required if SMTP alerts are enabled
SNMP	NMS (network management system) server	VIP, LAN1, LAN2, or MGMT	17 UDP	1024 → 65535	161	Required for SNMP management
SNMP Traps	MGMT or LAN1 on Grid Master or HA pair, or LAN1 on independent appliance	NMS server	17 UDP	1024 -> 65535	162	Required for SNMP trap management. Uses MGMT (when enabled) or LAN1 on Grid Master or HA pair, or LAN1 on independent appliance for the source address, depending on the destination IP address.

Service	SRC IP	DST IP	Proto	SRC Port	DST Port	Notes
SSHv2	Client	LAN1, LAN2, VIP, or MGMT on NIOS appliance	6 TCP	1024 -> 65535	22	Administrators can make an SSHv2 connection to the LAN1, LAN2, VIP, or MGMT port  Optional for management
Syslog	LAN1, LAN2, or MGMT of NIOS appliance	syslog server	17 UDP	1024 → 65535	514	Required for remote syslog logging
Traceroute	LAN1, LAN2, or UNIX-based appliance	VIP, LAN1, LAN2, or MGMT, or client	17 UDP	1024 → 65535	33000 → 65535	NIOS appliance responds with ICMP type code 3 (port unreachable)
TFTP Data	LAN1 or MGMT	TFTP server	17 UDP	1024 → 65535	69, then 1024 → 63999	For contacting a TFTP server during database and configuration backup and restore operations
VRRP	HA IP on the active node of HA pair	Multicast address 224.0.0.18	112 VRRP	802		For periodic announcements of the availability of the HA node that is linked to the VIP. The nodes in the HA pair must be in the same subnet.
HTTP	Management System	VIP, LAN1, or MGMT	6 TCP	1024 -> 65535	80	Required if the HTTP-redirect option is set on the Grid properties security page
HTTPS/SSL	Management System	VIP, LAN1, or MGMT	6 TCP	1024 → 65535	443	Required for administration through the GUI
Reporting	Reporting Forwarders	LAN1, LAN2, or MGMT on the indexer	6 TCP	1024 - 65535	9997	Required for the reporting service. Communication is single directional from forwarders to the indexer. For example, a forwarder detects events and forwards them to the indexer.
Reporting - Peer Replication	All Reporting Members	LAN1, LAN2, MGMT on each reporting member	TCP	1024 - 65535	7887	Splunk cluster peer replication (traffic among reporting members)

Service	SRC IP	DST IP	Proto	SRC Port	DST Port	Notes
Distributed Search	All Reporting Members	LAN1, LAN2, MGMT on each reporting member	TCP	1024 - 65535	7089	Distributed searches from Search Head to Reporting Members
Reporting Management	All Reporting Members	LAN1, LAN2, MGMT on each reporting member	TCP	1024 - 65535	8089	Grid Master to reporting members
Reporting Management	All Reporting Members	LAN1, LAN2, MGMT on each reporting member	TCP – IPv4	1024 - 65535	8000	Grid Master to reporting members
Reporting Management	All Reporting Members	LAN1, LAN2, MGMT on each reporting member	TCP – IPv6	1024 - 65535	8000	Grid Master to reporting members
Threat Protection	VIP on HA Grid Master or MGMT on single appliance (with threat protection service running)	N/A (using FQDN = <a href="https://ts.infoblox.com">https://ts.infoblox.com</a> )  This URL is configured to work with NIOS appliances. It has a self-signed certificate; it may not work properly with web browsers but works with appliances.	HTTPS	N/A	443	For threat protection rule updates.
Threat Insight	Client	N/A (using FQDN = <a href="https://ts.infoblox.com">https://ts.infoblox.com</a> )	HTTPS	N/A	443	For downloading module set and whitelist updates.
Microsoft Management	Managing Member	Microsoft Server	TCP	1024 - 65535	135, 445 Dynamic Port Range 1025-5000 (Windows Server 2003)  Dynamic Port Range 49152-65535 (Windows Server 2008)	Note that TCP ports 135 and 445 must be open on the Microsoft server, in addition to the dynamic port range. Ports 135 and 445 are used by the port mapper interface, which is a service on the Microsoft server that provides information to clients on which port to use to connect to a specific service, such as the service that allows the management of the DNS service.
DNS Forwarding to BloxOne Threat Defense Cloud: Cloud Services Portal	NIOS Appliance	BloxOne Threat Defense Cloud	TCP	443	443	csp.infoblox.com

Service	SRC IP	DST IP	Proto	SRC Port	DST Port	Notes
DNS Forwarding to BloxOne Threat Defense Cloud: Platform Management	NIOS Appliance	BloxOne Threat Defense Cloud	TCP	443	443	cp.noa.infoblox.com
DNS Forwarding to BloxOne Threat Defense Cloud: Application Management	NIOS Appliance	BloxOne Threat Defense Cloud	TCP	443	443	app.noa.infoblox.com
DNS Forwarding to BloxOne Threat Defense Cloud: NTP Server (Only if time sync with EXSi is disabled)	NIOS Appliance	BloxOne Threat Defense Cloud	UDP	123	123	ntp.ubuntu.com
DNS Forwarding to BloxOne Threat Defense Cloud: NTP Server (Only if time sync with EXSi is disabled)	NIOS Appliance	BloxOne Threat Defense Cloud	UDP	123	123	ubuntu.pool.ntp.org
DNS Forwarding to BloxOne Threat Defense Cloud: BloxOne Threat Defense Cloud DNS server	NIOS Appliance	BloxOne Threat Defense Cloud	UDP	123	123	52.119.40.100
SAML Authentication service	LAN1 or MGMT on Grid Master VIP on HA Grid Master		TCP	8765	Ports 443 (HTTPS) and 80 (HTTP)	

## Modifying Ethernet Port Settings

By default, the NIOS appliance automatically negotiates the optimal connection speed and transmission type (full or half duplex) on the physical links between the 10/100Base-T and 10/100/1000Base-T ports on the NIOS appliance and the Ethernet ports on a connecting switch. It is usually unnecessary to change the default auto-negotiation setting; however, you can manually configure connection settings for a port if necessary.

Occasionally, for example, even though both the NIOS appliance and the connecting switch support 1000-Mbps (megabits per second) full-duplex connections, they might fail to auto-negotiate that speed and type, and instead connect at lower speeds of either 100 or 10 Mbps using potentially mismatched full- and half-duplex transmissions. If this occurs, first determine if there is a firmware upgrade available for the switch. If so, apply the firmware upgrade and test the connection. If that does not resolve the issue, manually set the ports on the NIOS appliance and on the switch to make 1000-Mbps full-duplex connections.

To change Ethernet port settings:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.  
**Note:** You must enable the MGMT port before modifying its port settings. See [Using the MGMT Port](#).
2. In the **Network** tab of the *Grid Member Properties* editor, the Required Ports and Addresses table lists the network settings that were configured. This table lists the network settings of LAN1(IPv4) interface for an IPv4 member and LAN1(IPv6) interface for an IPv6 member. For a dual mode Grid member, this table lists the settings for both LAN1(IPv4) and LAN1(IPv6) interfaces. Complete the following to modify port settings:
  - **Interface:** Displays the name of the interface. You cannot modify this.
  - **Address:** Click the field and modify the IP address for the LAN1 port, which must be in a different subnet from that of the LAN2 and HA ports.
  - **Subnet Mask (IPv4) or Prefix Length (IPv6):** For IPv4 address, click the field and specify an appropriate subnet mask and for IPv6 address, specify the prefix length.
  - **Gateway:** Click the field and modify the default gateway for the LAN1 port.
  - **VLAN Tag:** Click the field and enter the VLAN tag ID if the port is configured for VLANs. You can enter a number from 1 to 4095.
  - **Port Settings:** From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
  - **DSCP Value:** Displays the Grid DSCP value. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

The port settings on the connecting switch must be identical to those you set on the NIOS appliance.

## Using the LAN2 Port



### Note

NIC failover for LAN1 and LAN2 is not supported on AWS members.

The LAN2 port is a 10/100/1000Base-T Ethernet connector on the front panel of the TE-815, TE-825, TE-1415, TE-1425, TE-2215, TE-2225, TE-4015 and TE-4025 appliances. By default, the LAN2 port is disabled and the appliance uses the LAN1 port (and HA port when deployed in an HA pair). Before you can enable and configure the LAN2 port on a Grid member, you must first configure the member and join it to the Grid. You must also have read/write permission to the Grid member on which you want to enable the port. When you enable the LAN2 port and SNMP, the appliance sends traps from this port for LAN2 related events.

You can configure the LAN2 port in different ways. You can enable the port redundancy or port failover feature, which groups the LAN1 and LAN2 ports into one logical interface. The LAN1/LAN2 grouping can be activated for both IPv4 and IPv6. Alternatively, you can configure the LAN2 port on a different IP network than LAN1, and enable the LAN2 port to provide DNS and DHCP services.

Note that you cannot use the LAN2 port to access the GUI and the API, or to connect to the Grid. This can impact the ability of other appliances, such as the NetMRI and PortIQ appliances, to communicate with the Grid Master. Any IPv6 services enabled for the LAN2 port also require provisioning of an IP address on the LAN2 port.

## About Port Redundancy

You can configure the LAN2 or LAN2 (VLAN) port to provide redundancy and additional fault tolerance in your network. Port redundancy is transparently supported for both IPv4 and IPv6. When you enable port redundancy, the LAN1 or LAN1 (VLAN) and LAN2 or LAN2 (VLAN) ports are grouped into one logical interface. They share one IP address and appear as one interface to the network. Then, if a link to one of the ports fails or is disabled, the appliance fails over to the other port, avoiding a service disruption.



You can connect the LAN1 or LAN1 (VLAN) and LAN2 or LAN2 (VLAN) ports to the same switch or to different switches, but the VLAN configuration between LAN1 and LAN2 must match. One port is active and the other port is idle at all times. In case of failure in the LAN1 or LAN1 (VLAN) port, the LAN2 or LAN2 (VLAN) port becomes active and once the LAN1 or LAN1 (VLAN) port is active again, the LAN2 or LAN2 (VLAN) port becomes passive if the **Use LAN1 when available** option is enabled. For more information, see [Enabling Port Redundancy](#) below.



#### Note

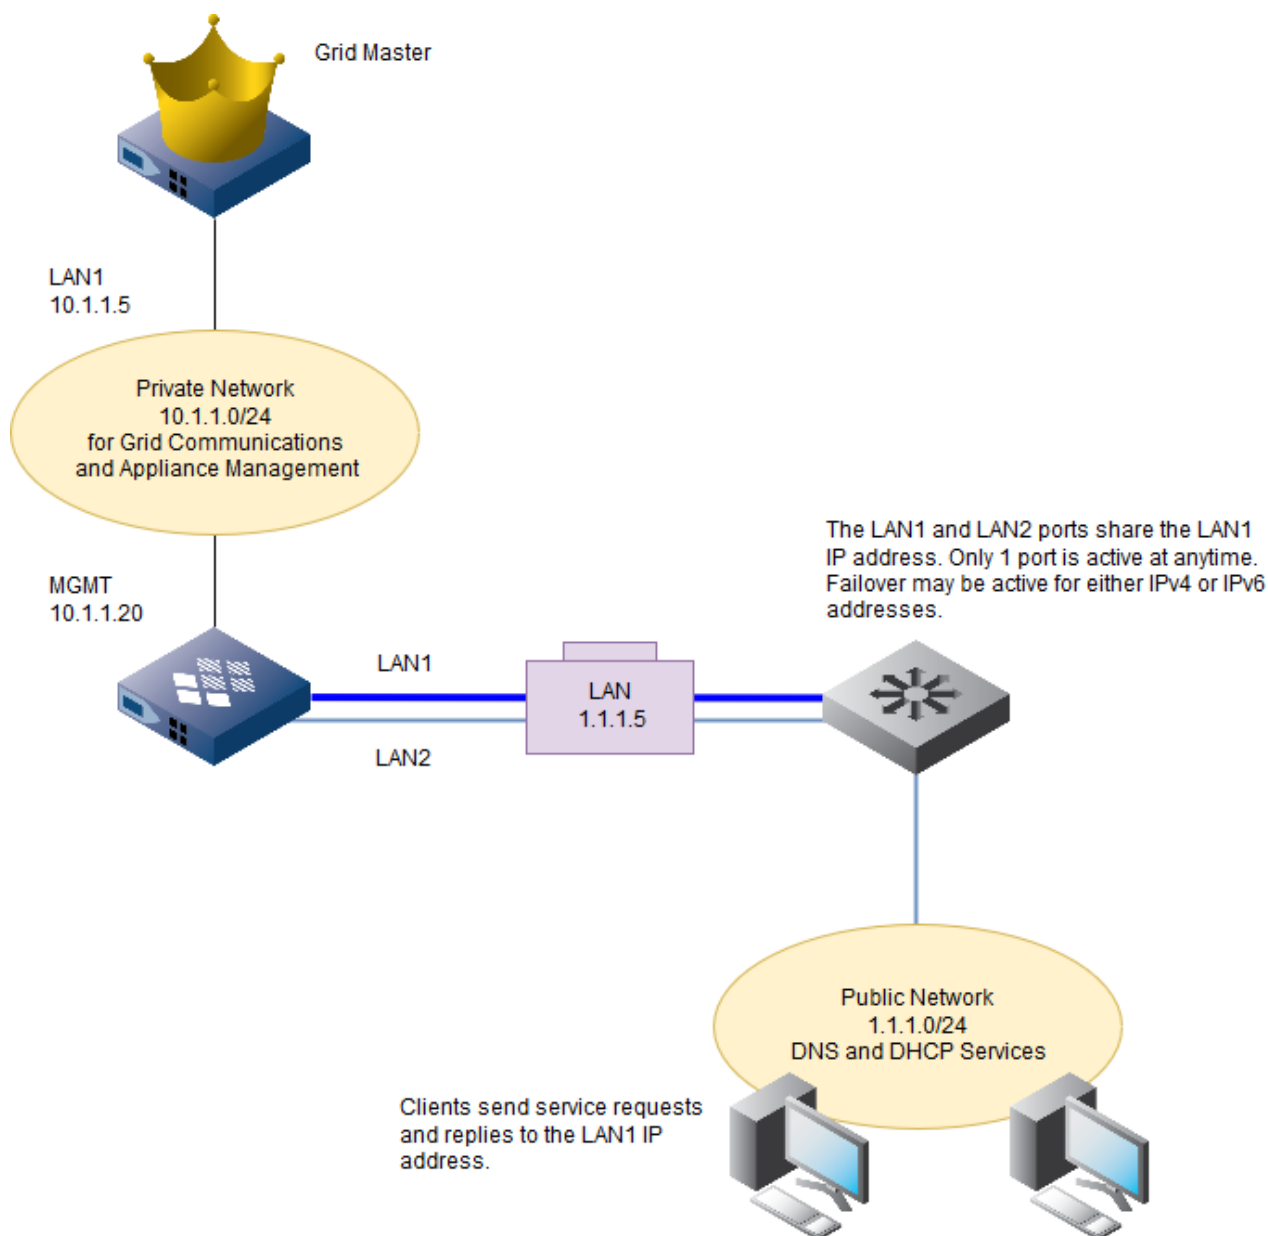
- When configuring port redundancy, the speed of the interfaces is not taken into consideration when selecting the active interface.

The LAN1 or LAN1 (VLAN) and LAN2 or LAN2 (VLAN) ports share the IP address of the LAN1 or LAN1 (VLAN) port; the port that is currently active owns the IP address. When you enable services on the appliance, such as DNS and DHCP, clients send their service requests to the LAN1 or LAN1 (VLAN) port IP address and receive replies from it as well. The port supports the services and features supported on the LAN1 or LAN1 (VLAN) port as listed in [Appliance Roles](#) table and [Sources and Destinations for Services](#) table, see [Configuring Ethernet Ports](#). You cannot enable the port redundancy feature if the LAN2 or LAN2 (VLAN) port is serving DNS or DHCP.

For example, you can use the MGMT port for Grid communications, and the LAN1 and LAN2 ports are connected to the same switch. The LAN1 and LAN2 port share the IP address of the LAN1 port, which is 1.1.1.5. In the illustration, LAN1 is the active port.

You can also have the MGMT port disabled and configure LAN1 and LAN2 for port redundancy. You can enable port redundancy on single or HA independent appliances and Grid members.

*Using the LAN2 Failover Feature*



## Enabling Port Redundancy

Before you enable port redundancy, ensure that both LAN1 and LAN2 are enabled. To enable port redundancy:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. In the **Network** -> **Basic** tab of the *Grid Member Properties* editor, select the **Enable port redundancy on LAN1/LAN2** checkbox.
3. Select the **Use LAN1 when available** checkbox to enable the NIOS appliance to automatically revert to the LAN1 port when it is available. In case of a connection failure, if the LAN1 port is not available, the NIOS appliance fails over to the LAN2 port. If you do not enable this option, the NIOS appliance will not automatically revert from the LAN2 port to the LAN1 port when it is available.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

The *Detailed Status* panel displays the status of both the LAN1 and LAN2 ports. In an HA pair, both nodes display the port information when port redundancy is enabled.

## Configuring the LAN2 Port

Before you enable the LAN2 port to provide DHCP and DNS services, you must specify its IP address and other properties. You can configure both IPv4 and IPv6 addresses for the LAN2 port of an IPv4, IPv6 and dual mode (IPv4 and IPv6) Grid member.

To configure the LAN2 port:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. In the **Network** -> **Basic** tab of the *Grid Member Properties* editor, click the Add icon of the **Additional Ports and Addresses** table and select **LAN2 (IPv4)** or **LAN2 (IPv6)** from the drop-down list. Enter the following:
  - **Interface:** Displays the name of the interface. You cannot modify this.
  - **Address:** Type the IP address for the LAN2 port, which must be in a different subnet from that of the LAN1 and HA ports.
  - **Subnet Mask (IPv4) or Prefix Length (IPv6):** Specify an appropriate subnet mask for IPv4 address and prefix length for IPv6 address.
  - **Gateway:** Type the default gateway for the LAN2 port.
  - **VLAN Tag:** Enter the VLAN tag ID if the port is configured for VLANs. You can enter a number from 1 to 4095. For information about VLAN, see [About Virtual LANs](#).
  - **Port Settings:** From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
  - **DSCP Value:** Displays the Grid DSCP value. To modify, click **Override** and then enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP and Implementing Quality of Service Using DSCP, see [Configuring Ethernet Ports](#).
  - **LAN2 Virtual Router ID (if HA):** If the appliance is in an HA pair, enter a VRID number.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

The *Detailed Status* panel displays the status of the LAN2 port. In an HA pair, only the active node displays the LAN2 information.

## Configuring LAN1/LAN2 for Automated Failover

You can use both LAN1 and LAN2 interfaces for DNS recursion. Both these interfaces have different gateways and you can send the DNS query source using the ANY IP address. If the default route interface goes down, the route redundancy feature configures another working interface so that there is automatic failover of recursion traffic from the failed interface to the working interface. This provides for a seamless flow of recursive traffic movement.

To configure automated failover, ensure that LAN1 and LAN2 have the same network configuration. Automated failover is supported on VLAN, IPv4, and IPv6 configurations.

To enable route redundancy:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. In the **Network** -> **Basic** tab of the *Grid Member Properties* editor, select the **Enable default route redundancy on LAN1/LAN2** checkbox.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

The *Detailed Status* panel displays the status of both the LAN1 and LAN2 ports.

If you select the **Enable default route redundancy on LAN1/LAN2** checkbox and then run the `show routes` CLI command, the output displays two default routes each having a different metric number. The primary default route that is set using the `set default_route LAN1|LAN2` CLI command has a metric value of 0 and the secondary default route has a metric value of 10 for IPv4 networks. For IPv6 networks, the primary default route has a metric value of 1024 and the secondary default route has a metric value of 1124. For more information, see [show routes](#).

In case of a failover, there may be a delay of a few seconds before the switchover to the secondary interface occurs.

If you want to enable port redundancy on LAN2 using WAPI, you have to use the entire `lan2_port_setting` struct even though the LAN2 IP addresses are already configured. The `lan2_port_setting` struct does not support partial updates.

## Enabling DHCP on LAN2

You can configure an appliance to provide DHCP service through the LAN1 port, LAN2 port, or both the LAN1 and LAN2 ports. Note that when you enable both ports, they must be connected to different subnets. You can also start and stop DHCP service for IPv4 or IPv6 on the LAN1 or LAN2 port after you have enabled the service.

After you configure the LAN2 port, you can enable DHCP services on the LAN2 port as follows:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. *If you are running DHCP for IPv4:* In the **General** -> **Basic** tab of the *Member DHCP Configuration* editor, select the **IPv4** checkbox for **LAN2** under DHCP Interfaces.  
*If you are running DHCP for IPv6:* In the **General** -> **Basic** tab of the *Member DHCP Configuration* editor, select the **IPv6** checkbox for **LAN2** under DHCP Interfaces. (An IPv6 address must also be provisioned for the port.) You can run either or both protocols for DHCP depending on your network deployment.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Enabling DNS on LAN2

If you enable DNS on an appliance, it always serves DNS on the LAN1 port. Optionally, you can configure the appliance to provide DNS services through the LAN2 port as well. For example, the appliance can provide DNS services through the LAN1 port for internal clients on a private network, and DNS services through the LAN2 port for external clients on a public network.

After you configure the LAN2 port, you can enable DNS services on the LAN2 port as follows:

1. From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. In the **General** -> **Basic** tab of the *Member DNS Configuration* editor, do the following:  
*If you are running DNS for IPv4:* In the **General** -> **Basic** tab of the *Member DHCP Configuration* editor, select the **IPv4** checkbox for **LAN2** under DNS Interfaces.  
*If you are running DNS for IPv6:* In the **General** -> **Basic** tab of the *Member DHCP Configuration* editor, select the **IPv6** checkbox for **LAN2** under DNS Interfaces. (An IPv6 address must also be provisioned for the port.) You can run either or both protocols for DNS depending on your network deployment.
  - **Automatically create glue A and PTR records for LAN2's address:** The NIOS appliance can automatically generate A (address) and PTR records for a primary name server whose host name belongs to the name space of the zone. Select this checkbox to enable the appliance to automatically generate an A and PTR record.
  - **Automatically create IPv6 glue AAAA and PTR records for LAN2's address:** automatically generate AAAA and PTR records for the LAN2 IPv6 address. A glue record is the IP address of a name server held at the domain name registry. They are needed to set a domain's name server to a host name within the domain. Example: to set the name servers of ns1.corpxyz.com and ns2.corpxyz.com, provide the glue records, which are in effect the IP addresses, for ns1.corpxyz.com and ns2.corpxyz.com, within specific DNS record types.  
Without the glue records, DNS requests never resolve to the correct IP address because the domain registry does not associate the IP with the correct records.
3. In the **General** -> **Advanced** tab (click **Toggle Advanced Mode** if necessary), select one of the following from the **Send queries from** and the **Send notify messages and zone transfer request from** drop-down lists:
  - **VIP:** The appliance uses the IP address of the HA port as the source for queries, notifies, and zone transfer requests.
  - **MGMT:** The appliance uses the IP address of the MGMT port as the source for queries, notifies, and zone transfer requests.

- **LAN2:** The appliance uses the IP address of the LAN2 port as the source for queries, notifies, and zone transfer requests.
  - **Any:** The appliance chooses which port to use as the source for queries, notifies, and zone transfer requests.  
The **Send queries from** drop-down list also includes loopback IP addresses that you configured. You can select a loopback address as the source for queries.
4. Save the configuration and click **Restart** if it appears at the top of the screen.
  5. Click **Restart** to restart services.

## Using the MGMT Port

The MGMT (Management) port is a 10/100/1000Base-T Ethernet connector on the front panel of the TE-815, TE-825, TE-1415, TE-1425, TE-2215, TE-2225, TE-4015 and TE-4025 appliances. It allows you to isolate the following types of traffic from other types of traffic on the LAN and HA ports:

- Appliance Management
- Grid Communications
- DNS Services

For information about what types of traffic qualify as appliance management, Grid communications, and DNS services, see the Sources and Destinations for Services table in [Configuring Ethernet Ports](#).



### Note


The MGMT port currently does not support DHCP, NAT, or TFTP. IPv6 addressing may be applied to the MGMT port.


Some NIOS appliance deployment scenarios support more than one concurrent use of the MGMT port. The following table depicts MGMT port uses for various appliance configurations.

*Supported MGMT Port Uses for Various appliance Configurations*

Appliance Configuration	Appliance Management	Grid Communications	DNS Services
Single Independent Appliance		Not Applicable	
Independent HA Pair		Not Applicable	
Grid Master			
Grid Master Candidate			
HA Grid Member			
Single Grid Member			

\* Although you manage all Grid members through the Grid Master, if you enable the MGMT port on common Grid members, they can send syslog events, SNMP traps, and e-mail notifications, and receive SSH connections on that port.

Infoblox does not support MGMT port usage for some appliance configurations (indicated by the symbol  in Supported MGMT Port Uses for Various appliance Configurations table ) because it cannot provide redundancy through the use of a VIP. A Grid Master that is an HA pair needs the redundancy that a VIP interface on the HA port provides for Grid communications. Similarly, DNS servers in an HA pair need that redundancy to answer DNS queries. Because the MGMT port does not support a VIP and thus cannot provide redundancy, Grid Masters (and potential Grid Masters) do not support Grid communications on the MGMT port.

In addition, NIOS appliances in an HA pair support DNS services on the active node only (indicated by the symbol  in Supported MGMT Port Uses for Various appliance Configurations table ). Only the active node can respond to queries that it receives. If a DNS client sends a query to the MGMT port of the node that happens to be the passive node, the query can eventually time out and fail.

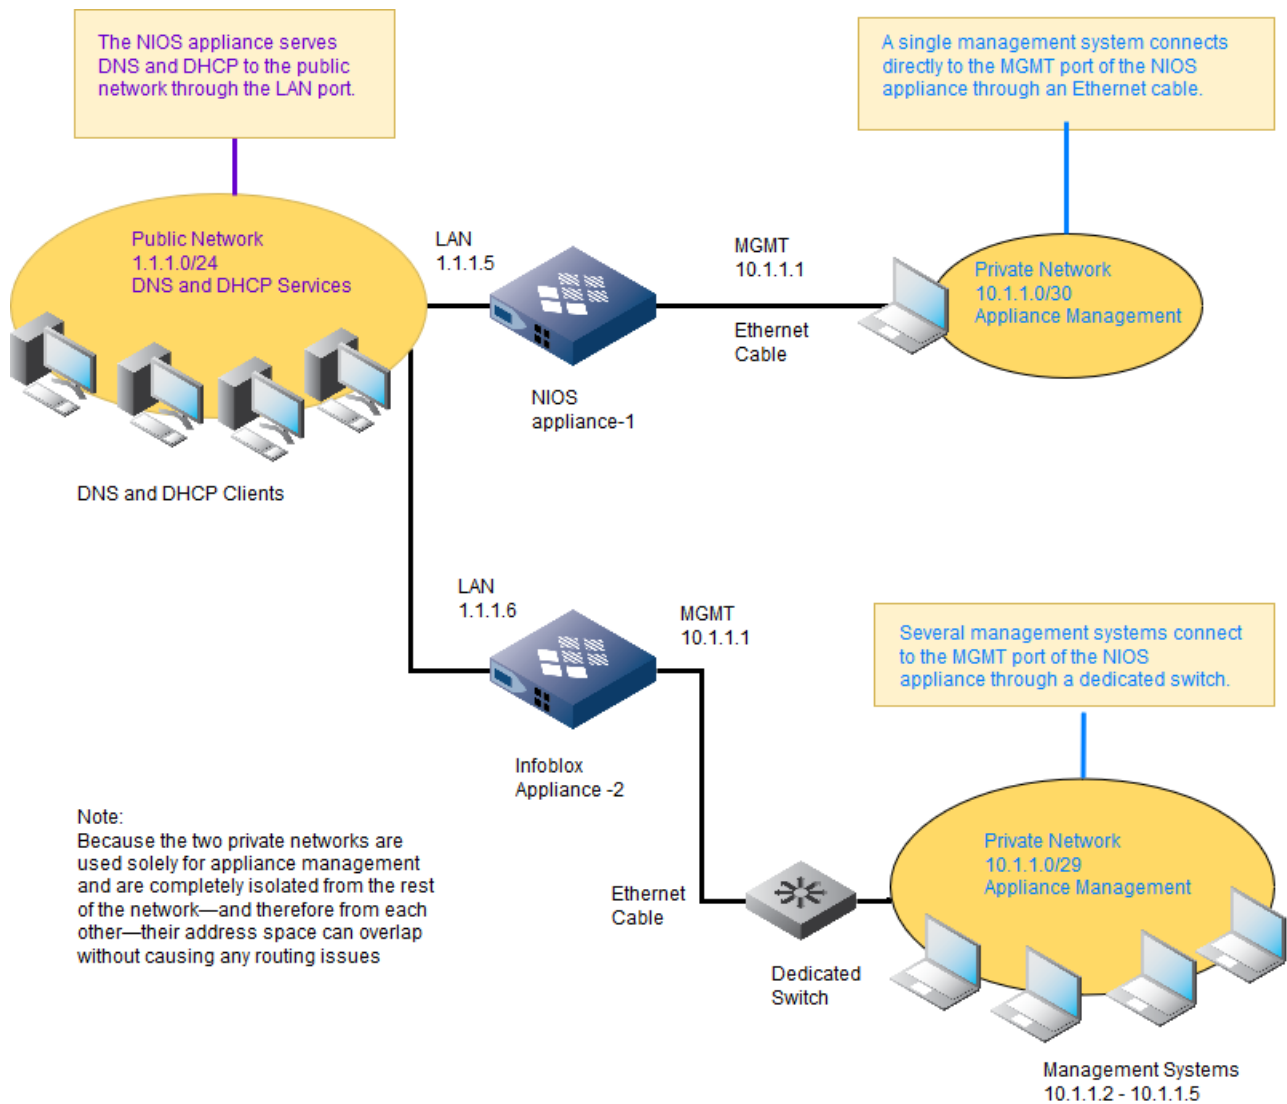
The MGMT port is not enabled by default. By default, a NIOS appliance uses the LAN port (and HA port when deployed in an HA pair). You must log in using a superuser account to enable and configure the MGMT port. You can configure both IPv4 address and IPv6 address for the MGMT port of a Grid member. You can enable the MGMT port through the Infoblox GUI, as explained in the following sections.

## Appliance Management

You can restrict administrative access to a NIOS appliance by connecting the MGMT port to a subnet containing only management systems. This approach ensures that only appliances on that subnet can access the Infoblox GUI and receive appliance management communications such as syslog events, SNMP traps, and e-mail notifications from the appliance.

If you are the only administrator, you can connect your management system directly to the MGMT port. If there are several administrators, you can define a small subnet—such as 10.1.1.0/29, which provides six host IP addresses (10.1.1.1–10.1.1.6) plus the network address 10.1.1.0 and the broadcast address 10.1.1.7—and connect to the NIOS appliance through a dedicated switch (which is not connected to the rest of the network). The figure below shows how an independent appliance separates appliance management traffic from network protocol services. Note that the LAN port is on a different subnet from the MGMT port.

*Appliance Management from One or More Management Systems*



Similarly, you can restrict management access to a Grid Master to only those appliances connected to the MGMT ports of the active and passive nodes of the Grid Master.

To enable the MGMT port on an independent appliance or Grid Master for appliance management and then cable the MGMT port directly to your management system or to a network forwarding appliance such as a switch or router:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. In the **Network** -> **Basic** tab of the *Grid Member Properties* editor, add the MGMT port to the Additional Ports and Addresses table as follows:
3. Click the Add icon and select **MGMT (IPv4)** to configure an IPv4 address or select **MGMT (IPv6)** to configure an IPv6 address for the MGMT port. You can configure both IPv4 and IPv6 addresses for the MGMT port. Grid Manager adds a row for the MGMT port. For an HA pair, it adds two rows, one for each node.
4. Enter the following in the row of the MGMT port for a single Grid Master or independent appliance, and in the rows of the two nodes for an HA Grid Master or independent HA pair:
  - **Interface:** Displays the name of the interface. You cannot modify this.
  - **Address:** Type the IP address for the MGMT port, which must be in a different subnet from that of the LAN and HA ports.
  - **Subnet Mask (IPv4) or Prefix Length (IPv6):** For IPv4 address, specify an appropriate subnet mask for the number of management systems that you want to access the appliance through the MGMT port. For IPv6 address, specify the prefix length.

- **Gateway:** Type the default gateway for the MGMT port. If you need to define any static routes for traffic originating from the MGMT port—such as SNMP traps, syslog events, and email notifications—destined for remote subnets beyond the immediate subnet, specify the IP address of this gateway in the route.
- **Port Settings:** Choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
- **DSCP Value:** Displays the Grid DSCP value. To modify, click **Override** and then enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP and Implementing Quality of Service Using DSCP, see [Configuring Ethernet Ports](#).

1. In the **Network** -> **Advanced** tab, make sure that the **Enable VPN on MGMT Port** checkbox is not selected.
2. Save the configuration and click **Restart** if it appears at the top of the screen.
3. Log out of Grid Manager.
4. Cable the MGMT port to your management system or to a switch or router to which your management system can also connect.
5. If your management system is in a subnet from which it cannot reach the MGMT port, move it to a subnet from which it can.  
The Infoblox Grid Manager GUI is now accessible through the MGMT port on the NIOS appliance from your management system.
6. Open an Internet browser window and enter the IP address of the MGMT port as follows: *https://<IP address of MGMT port>*.
7. Log in to Grid Manager.
8. Check the *Detailed Status* panel of the Grid member to make sure the status icons are green.

## Grid Communications

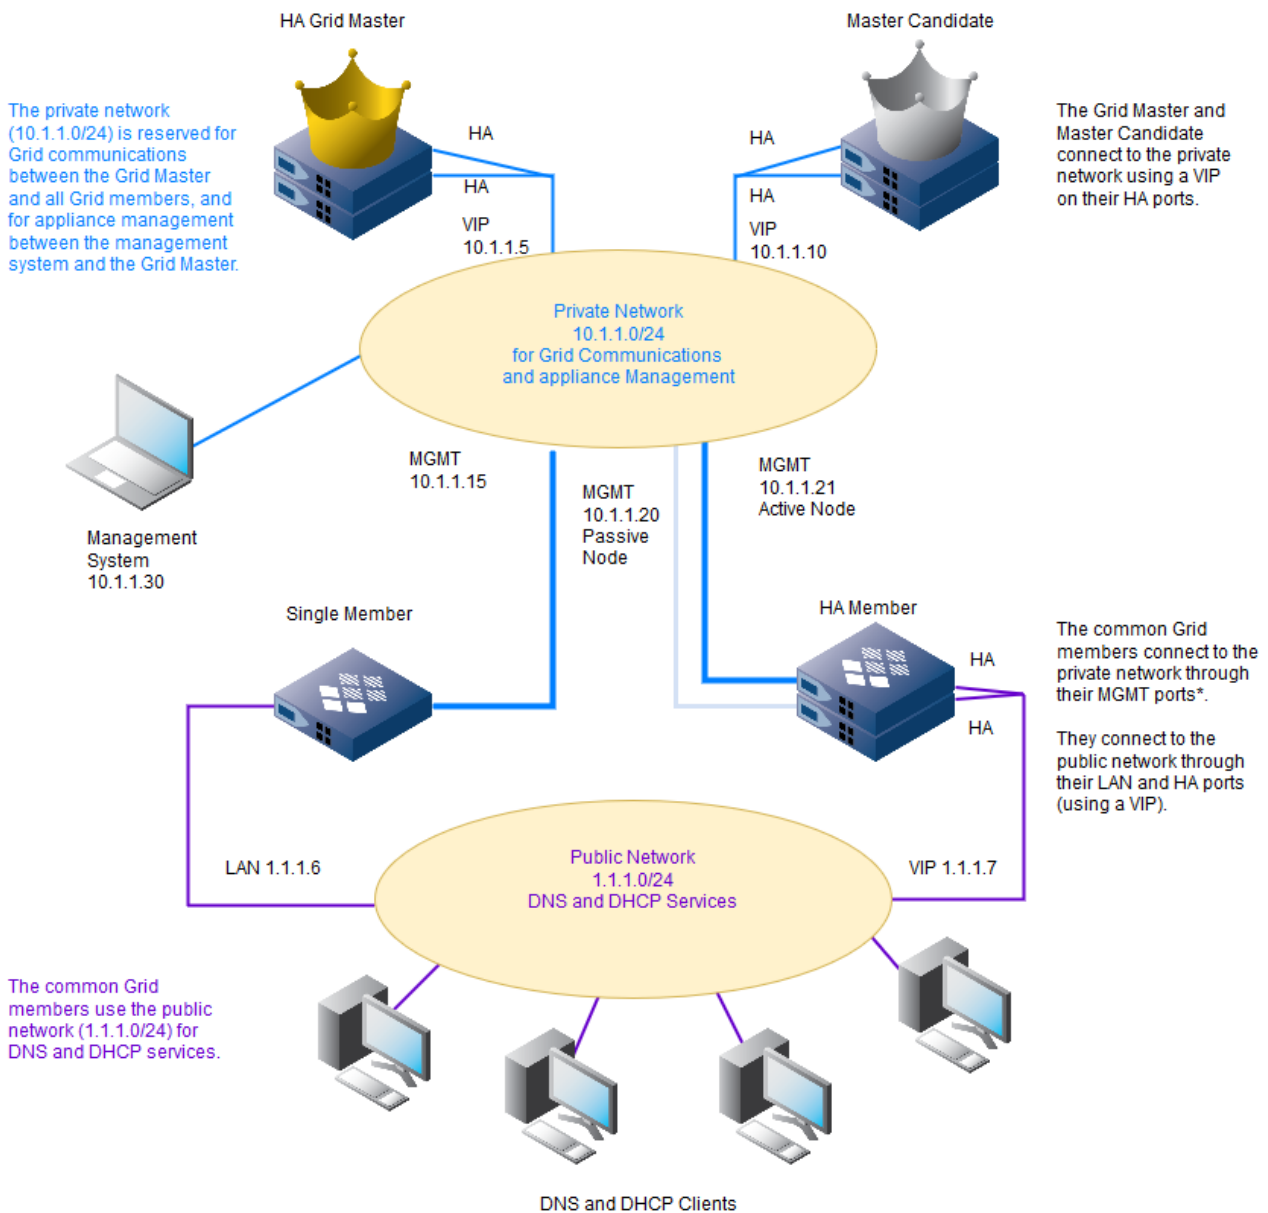
You can isolate all Grid communications to a dedicated subnet as follows:

- For Grid communications from the Grid Master, which can be an HA pair or a single appliance, the master uses either the VIP interface on the HA port of its active node (HA master) or its LAN port (single master). Neither a single nor HA Grid Master can use its MGMT port for Grid communications. (This restriction applies equally to Master Candidates.)
- Common Grid members connect to the Grid Master through their MGMT port.

This ensures that all database synchronization and Grid maintenance operations are inaccessible from other network elements while the common Grid members provide network protocol services on their LAN ports. The below figure shows how Grid members communicate to the master over a dedicated subnet.

### *Grid Communications*





\* Only the active node of an HA member connects to the Grid Master. The passive node communicates just with the active node. If there is an HA failover, the newly promoted active node must first join the Grid before continuing Grid communications with the Grid Master on behalf of the HA member.

## Enabling Grid Communications over the MGMT Port for Existing Grid Members

To enable the MGMT port for Grid communications on an existing single or HA Grid member:

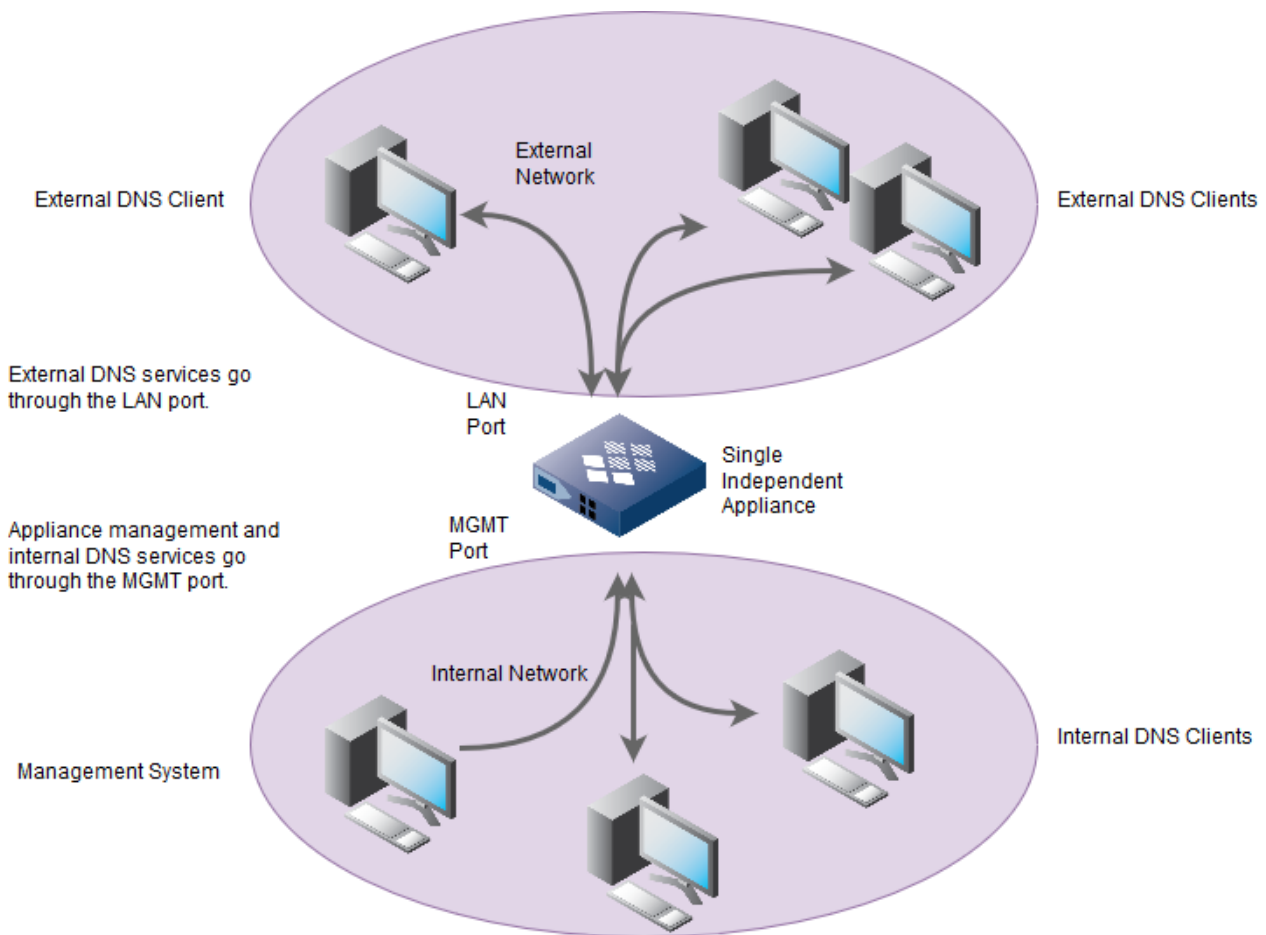
1. Log in to the Grid Master with a superuser account.
2. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.  
Note that you must enable the MGMT port before modifying its port settings. See the Using the MGMT Port page.
3. In the **Network** -> **Basic** tab of the *Grid Member Properties* editor, add the MGMT port to the Additional Ports and Addresses table as follows:

4. Click the Add icon and select **MGMT (IPv4)** to configure an IPv4 address or select **MGMT (IPv6)** to configure an IPv6 address for the MGMT port. You can configure both IPv4 address and IPv6 address for the MGMT port. Grid Manager adds a row for the MGMT port. For an HA pair, it adds two rows, one for each node.
5. Enter the following in the row of the MGMT port for a single Grid Master or independent appliance, and in the rows of the two nodes for an HA Grid Master or independent HA pair:
  - **Interface:** Displays the name of the interface. You cannot modify this.
  - **Address:** Type the IP address for the MGMT port, which must be in a different subnet from that of the LAN and HA ports.
  - **Subnet Mask (IPv4) or Prefix Length (IPv6):** For IPv4 address, specify an appropriate subnet mask for the number of management systems that you want to access the appliance through the MGMT port. For IPv6 address, specify the prefix length.
  - **Gateway:** Type the default gateway for the MGMT port. If you need to define any static routes for traffic originating from the MGMT port—such as SNMP traps, syslog events, and email notifications—destined for remote subnets beyond the immediate subnet, specify the IP address of this gateway in the route.
  - **Port Settings:** Choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
  - **DSCP Value:** Displays the Grid DSCP value. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP and Implementing Quality of Service Using DSCP, see [Configuring Ethernet Ports](#).
6. In the **Network** -> **Advanced** tab, select the **Enable VPN on MGMT Port** checkbox.
7. In the **Security** tab, do the following:
  - **Restrict Remote Console and Support Access to MGMT Port:** Select this checkbox to restrict SSH (Secure Shell) v2 access to the MGMT port only. This restricts Infoblox Technical Support and remote console connections—both of which use SSH v2—to just the MGMT port. For an HA pair, you can make an SSH v2 connection to the MGMT port on both the active and passive nodes. Clear the checkbox to allow SSH v2 access to both the MGMT and LAN ports. For an HA pair, you can make an SSH v2 connection to the MGMT and LAN ports on both the active and passive nodes.
8. Save the configuration and click **Restart** if it appears at the top of the screen. The master communicates the new port settings to the member, which immediately begins using them. The member stops using its LAN port for Grid communications and begins using the MGMT port.
9. To confirm that the member still has Grid connectivity, check that the status icons for that member are green on the *Detailed Status* and *Grid* panels.

## DNS Services

You can configure a single independent appliance or single Grid member to provide DNS services through the MGMT port in addition to the LAN port. For example, the appliance can provide DNS services through the MGMT port for internal clients on a private network, and DNS services through the LAN port for external clients on a public network. While providing DNS services on the MGMT port, you can still use that port simultaneously for appliance management. The figure below shows a management system communicating with a single independent appliance through its MGMT port while the appliance also provides DNS services on that port to a private network. Additionally, the appliance provides DNS services to an external network through its LAN port.

*DNS Services on the LAN and MGMT Ports, and appliance Management on the MGMT Port*



Like a single independent appliance, a single Grid member can also support concurrent DNS traffic on its MGMT and LAN ports. However, because you manage all Grid members through the Grid Master, a Grid member only uses an enabled MGMT port to send SNMP traps, syslog events, and email notifications, and to receive SSH connections. In addition, the active node of an HA pair can provide DNS services through its MGMT port. To use this feature, you must enable DNS services on the MGMT ports of both nodes in the HA pair and specify the MGMT port IP addresses of both nodes on the DNS client as well, in case there is a failover and the passive node becomes active. Note that only the active node can respond to queries that it receives. If a DNS client sends a query to the MGMT port of the node that happens to be the passive node, the query can eventually time out and fail. To enable DNS services on the MGMT port of an appliance:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.  
**Note:**  
 You must enable the MGMT port before modifying its port settings.  
 You must mandatorily configure the LAN interface before joining the HA nodes to the Grid. If you join the nodes with VLAN tagging already enabled on HA, the new nodes must join with VLAN tagging only. If you join the nodes using the MGMT interface, you must enable VLAN tagging for the new nodes.
2. In the **Network** -> **Basic** tab of the *Grid Member Properties* editor, add the MGMT port to the Additional Ports and Addresses table as follows:
3. Click the Add icon and select **MGMT (IPv4)** to configure an IPv4 address or select **MGMT (IPv6)** to configure an IPv6 address for the MGMT port. You can configure both IPv4 and IPv6 address for the MGMT port. Grid Manager adds a row for the MGMT port. For an HA pair, it adds two rows, one for each node.
4. Enter the following in the row of the MGMT port for a single Grid Master or independent appliance, and in the rows of the two nodes for an HA Grid Master or independent HA pair:
  - **Interface:** Displays the name of the interface. You cannot modify this.

- **Address:** Type the IP address for the MGMT port, which must be in a different subnet from that of the LAN and HA ports.
  - **Subnet Mask (IPv4) or Prefix Length (IPv6):** For IPv4 address, specify an appropriate subnet mask for the number of management systems that you want to access the appliance through the MGMT port. For IPv6 address, specify the prefix length.
  - **Gateway:** Type the default gateway for the MGMT port. If you need to define any static routes for traffic originating from the MGMT port—such as SNMP traps, syslog events, and email notifications—destined for remote subnets beyond the immediate subnet, specify the IP address of this gateway in the route.
  - **Port Settings:** Choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.
  - **DSCP Value:** Displays the Grid DSCP value. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP and Implementing Quality of Service Using DSCP, see [Configuring Ethernet Ports](#).
5. Click **Save & Close** to save your settings for the MGMT port.
  6. From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
  7. In the **General** -> **Basic** tab of the *Member DNS Properties* editor, do the following:
    - If you are running DNS service for IPv4, select the IPv4 checkbox for **MGMT** under **DNS Interfaces**.
    - If you are running DNS service for IPv6, select the IPv6 checkbox for **MGMT** under **DNS Interfaces**.
  8. In the **General** -> **Advanced** tab, select one of the following from the **Send queries from** and the **Send notify messages and zone transfer requests from** drop-down lists:
    - **VIP:** The appliance uses the IP address of the HA port as the source for queries, notifies, and zone transfer requests.
    - **MGMT:** The appliance uses the IP address of the MGMT port as the source for queries, notifies, and zone transfer requests.
    - **LAN2:** The appliance uses the IP address of the LAN2 port as the source for queries, notifies, and zone transfer requests.
    - **Any:** The appliance chooses which port to use as the source for queries, notifies, and zone transfer requests.

The **Send queries from** drop-down list also includes loopback IP addresses that you configured. You can select a loopback address as the source for queries.
  9. Save the configuration and click **Restart** if it appears at the top of the screen.

To see that the appliance now also serves DNS on the MGMT port:

1. From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *Grid\_member* checkbox.
2. Expand the Toolbar and click **View** -> **View DNS Configuration**.
3. Check that the IP address of the MGMT port appears in the address match list in the listen-on substatement.

## About Lights Out Management

Infoblox LOM (Lights Out Management) is an implementation of the remote management and monitoring of Infoblox appliances that are LOM ready, such as the Trinzic 2010 appliances.

The LOM feature is useful when you want to monitor your platforms remotely or consolidate your data centers. When you monitor your systems remotely, you can avoid issues such as overheating of a problematic system by remotely powering down the system. To conserve energy, you can also power up and down any systems based on service requirements. You can enable LOM for the entire Grid and override the Grid settings for specific members. You can also configure LOM on independent appliances and HA pairs.



### Note

You can configure LOM only on appliances that support LOM. This port automatically negotiates a speed of up to 1000 Mbps. Devices connected to the LOM port must be configured to auto-negotiate and must not have a fixed speed of 1000 Mbps.

LOM is disabled by default. Before you can configure LOM and remotely manage the appliance, ensure that the IPMI (Intelligent Platform Management Interface) port on your appliance is properly connected to the network. Consider the following security measures before you enable the IPMI interface for LOM:

- Use an authentication method other than the RAKP (Remote Authenticated Key-Exchange Protocol) for the IPMI interface. Any implementation that uses the RAKP can become vulnerable.
- Secure the network to which the IPMI interface is connected. Infoblox recommends that you use a separate and secure network for all IPMI traffic.
- Use strong passwords for all IPMI users. At least 10 random characters are recommended. Attacks are only effective against weak passwords.
- IPMI is disabled by default. DO NOT enable IPMI on the appliances if it is not being used.

By default, IPMI uses UDP port 623. You can then enable LOM and add LOM users through the Infoblox GUI. When you add LOM users, you can assign them specific roles so they can perform only certain functions. When you add a LOM user, you can configure the user to be an "operator" or a "user" depending on the functions you want the user to perform. An operator can access an appliance remotely and perform the following functions:

- Access the serial console
- Reset the appliance
- Power up and down the appliance
- Monitor system status, such as CPU usage and system temperature

A user role can only monitor system status. Users with this role cannot perform any other functions remotely. After you set up and configure your appliance, perform the following tasks through Grid Manager to enable LOM and set up LOM users:

1. Enable LOM for the Grid or members that support IPMI, as described in [Enabling LOM](#) below.
2. Add LOM users based on your organizational needs, as described in [Adding LOM User Accounts](#) below.
3. Configure the IPMI network interface on the appliance, as described in [Configuring the IPMI Network Interface](#), see [Configuring SNMP](#).
4. After you have configured LOM and set up the IPMI interface, install a utility such as IPMITool on your Linux management system. For information about IPMITool, visit the IPMITool web site at <http://ipmitool.sourceforge.net>. For the most commonly used commands and examples, see [IPMI Commands and Examples](#) below.

You can also do the following from Grid Manager after you configure LOM:

- Enable and disable LOM for the Grid or members, as described in [Enabling LOM](#) below.
- Modify LOM settings, as described in [Modifying LOM Settings](#) below.
- View LOM users, as described in [Modifying LOM Settings](#) below.

## Enabling LOM

Before you can add LOM users and manage Infoblox appliances remotely, you must enable LOM. When LOM is configured for the entire Grid, all members inherit the Grid settings. You can also override the Grid settings for specific members. For an HA pair, you can configure LOM on the node that supports IPMI.

To enable and disable LOM:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.  
**Independent appliance:** From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties -> Edit**.  
**Member:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the **LOM** tab, complete the following:
  - **Enable Lights Out Management:** LOM is disabled by default. Select this checkbox to enable LOM. When LOM is enabled or disabled for the Grid, all members inherit the same setting.
3. Save the configuration.

## Adding LOM User Accounts

You can add up to eight LOM user accounts. Admins must use the configured user name and password to remotely log in to the appliance.

Note that when you add LOM user accounts at the Grid level, all members inherit them. You can configure user accounts specific to a member by overriding the Grid accounts. When you click **Override** to modify the inherited Grid accounts, the appliance creates copies of the Grid level user accounts and saves them at the member level. These are new accounts at the member level and do not affect the Grid accounts or any accounts configured on other Grid members. You can also reset member accounts to the Grid accounts by clicking **Inherit**. When you do that however, all changes you previously made to the member accounts are lost.

To add a LOM user account:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.  
**Independent appliance:** From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties -> Edit**.  
**Member:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the **LOM** tab, complete the following:
  - **User Accounts:** Click the Add icon and complete the following:
    - **Name:** Enter the name of the LOM user account.
    - **Password:** Enter the password for the LOM user account. Note that while the maximum length allowed for the password is 15 characters and the minimum length is 4 characters. If you are running NIOS on the IB-4000 or the ND-4000 platform, you must enter a minimum password length of 8 characters. Otherwise, an error message is displayed.
    - **Confirm Password:** Enter the password again.
    - **Role:** From the drop-down list, select the role for the LOM user account. **Operator** allows users to perform all supported LOM related functions. **User** allows admins to only monitor system sensors such as temperature and CPU usage.
    - **Disable:** Select this to deactivate the user account but keep a user profile.
    - Click **Add** to add the new user account.
3. Save the configuration.

## Configuring the IPMI Network Interface

You must configure the IPMI network interface before you can access the appliance remotely. To configure the IPMI network interface:

1. **Independent appliance:** From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties -> Edit**.  
**Member:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. In the **LOM** tab, complete the following in the Network Configuration table:
  - **Address:** Enter the IPMI interface address here.
  - **Subnet Mask:** Enter the subnet mask for the IPMI interface.
  - **Gateway:** Enter the gateway address for the IPMI interface.
3. Save the configuration.

## Modifying LOM Settings

To modify LOM settings:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.  
**Independent appliance:** From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties -> Edit**.  
**Member:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.

2. Modify the following:
  - **Enable Lights Out Management:** LOM is disabled by default. Select this checkbox to enable LOM. When you enable or disable this for the Grid, all members inherit the same setting.
  - **Network Configuration:** Click the fields in the table to modify the IPv4 address, subnet mask, and gateway address for the IPMI interface. For an HA pair, the appliance displays information only for the nodes that support IPMI. Enter the information for the following fields: **Address**, **Subnet Mask**, and **Gateway**. The **Node** and **LAN Address** fields are read-only, and you cannot modify them. The LAN address is the IPMI interface address.
  - **User Accounts:** Click the Add icon to add new LOM users. You can also select an existing LOM user and click the Edit icon to modify the user settings, as described in Adding LOM User Accounts above.
3. Save the configuration.

## Viewing LOM Users

To view information about LOM users:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties -> Edit**.  
**Independent appliance:** From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties -> Edit**.  
**Member:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.  
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the **LOM** tab, Grid Manager displays the following information for each LOM user:
  - **Name:** The name of the LOM user.
  - **Role:** The user role to which the LOM user was assigned. This can be **Operator** or **User**.
  - **Disabled:** Indicates whether the LOM user account is disabled or not. When a LOM user account is disabled, the user cannot access the appliance remotely.

## IPMI Commands and Examples

This section describes some of the most commonly used IPMITool commands and examples. For more information about the IPMI commands and usage, visit the IPMITool web site a <http://ipmitool.sourceforge.net>.

To use IPMI commands, complete the following:

1. Ensure that you have properly enabled and configured LOM and the IPMI network interface.
2. Install IPMITool on a Linux management system. For information, visit the IPMITool web site at <http://ipmitool.sourceforge.net>.

Access IPMITool and enter an IPMI command to perform a specific task. The appliance displays the corresponding output.

Following are some of the most commonly used IPMI commands and their sample outputs. Note that command outputs vary by appliances. The following sample commands were performed on a Trinzic 1415 appliance. All sample commands in this section use the following syntax:

```
ipmitool -H <LOMIPAddress> -U username -P password -L [OPERATOR/USER] -I
lanplus
<supported commands>
```



### Note

If you reset IB-2225 using LOM and the interval between power off and power on is lower than 10 minutes, a hardware watchdog timeout message may be displayed in the syslog and may reset the newly powered on system thus putting it in an unstable state. If the interval is longer, the watchdog timer is released before the system is up, and a restart will bring the system in a good state without the watchdog timeout messages.



## Command to be Used with Caution

```
power reset variant
```

---

Caution: Using this command has the same effect as pulling the power cord off the appliance.

---

## Checking Power Status with User Role

Command:

```
ipmitool -H 10.37.2.70 -U user -P infoblox -L USER -I lanplus power status
```

Command output:

```
Chassis Power is on
```

## Checking Various Sensors [Temperature, Voltage, FANS, Physical Security, Power supply, OEM] with User Role

Command:

```
ipmitool -H 10.37.2.70 -U user -P infoblox -L USER -I lanplus sensor
```

Command output:

```
System Temp | 23.000 | degrees C | ok | -9.000 | -7.000 | -5.000 | 75.000 |  
77.000 | 79.000  
CPU Temp | 0x0 | discrete | 0x0000 | na | na | na | na | na | na  
FAN 1 | 10390.000 | RPM | ok | 215.000 | 400.000 | 585.000 | 29260.000 |  
29815.000 |  
30370.000  
FAN 2 | na | RPM | na | na | na | na | na | na | na  
FAN 3 | 9835.000 | RPM | ok | 215.000 | 400.000 | 585.000 | 29260.000 |  
29815.000 |  
30370.000  
FAN 4 | 11870.000 | RPM | ok | 215.000 | 400.000 | 585.000 | 29260.000 |  
29815.000 |  
30370.000  
FAN 5 | 10390.000 | RPM | ok | 215.000 | 400.000 | 585.000 | 29260.000 |  
29815.000 |  
30370.000  
CPU Vcore | 0.832 | Volts | ok | 0.640 | 0.664 | 0.688 | 1.344 | 1.408 | 1.472  
+3.3VCC | 3.264 | Volts | ok | 2.816 | 2.880 | 2.944 | 3.584 | 3.648 | 3.712  
+12 V | 11.978 | Volts | ok | 10.494 | 10.600 | 10.706 | 13.091 | 13.197 |  
13.303
```



```

CPU DIMM | 1.528 | Volts | ok | 1.152 | 1.216 | 1.280 | 1.760 | 1.776 | 1.792
+5 V | 5.088 | Volts | ok | 4.096 | 4.320 | 4.576 | 5.344 | 5.600 | 5.632
-12 V | -12.486 | Volts | ok | -13.844 | -13.650 | -13.456 | -10.934 | -10.740
| -10.546
VBAT | 3.120 | Volts | ok | 2.816 | 2.880 | 2.944 | 3.584 | 3.648 | 3.712
+3.3VSB | 3.264 | Volts | ok | 2.816 | 2.880 | 2.944 | 3.584 | 3.648 | 3.712
AVCC | 3.264 | Volts | ok | 2.816 | 2.880 | 2.944 | 3.584 | 3.648 | 3.712
Chassis Intru | 0x0 | discrete | 0x0000 | na | na | na | na | na | na PS Status
| 0x1 | discrete | 0x01ff | na | na | na | na | na | na

```

#### Printing System Event Log with User Role

Command:

```
ipmitool -H 10.37.2.70 -U user -P infoblox -L USER -I lanplus sel list
```

Command output: The appliance displays all event log entries (if any)

#### Getting FRU Information with User Role

Command:

```
ipmitool -H 10.37.2.70 -U user -P infoblox -L USER -I lanplus fru
```

Command output:

```

FRU Device Description : Builtin FRU Device (ID 0) Board Mfg Date : Sun Dec 31
15:00:00 1995
Board Mfg : Supermicro Board Serial : Product Serial :

```

#### Powering Off the Appliance with Operator Role

Command:

```
ipmitool -H 10.37.2.70 -U operator -P infoblox -L OPERATOR -I lanplus power off
```

Command output:

```
Chassis Power Control: Down/Off
```

#### Powering On the Appliance with Operator Role

Command:

```
ipmitool -H 10.37.2.70 -U operator -P infoblox -L OPERATOR -I lanplus power on
```

Command output:

```
Chassis Power Control: Up/On
```

#### Activating the Serial Console Port using Operator role

Command:

```
ipmitool -H 10.37.2.70 -U operator -P infoblox -L OPERATOR -I lanplus sol
```

```
activate
```

Command output:

```
[SOL Session operational. Use ~? for help] login: admin
```

```
password:
```

```
Infoblox NIOS Release 6.4.0-163715 (64bit)
```

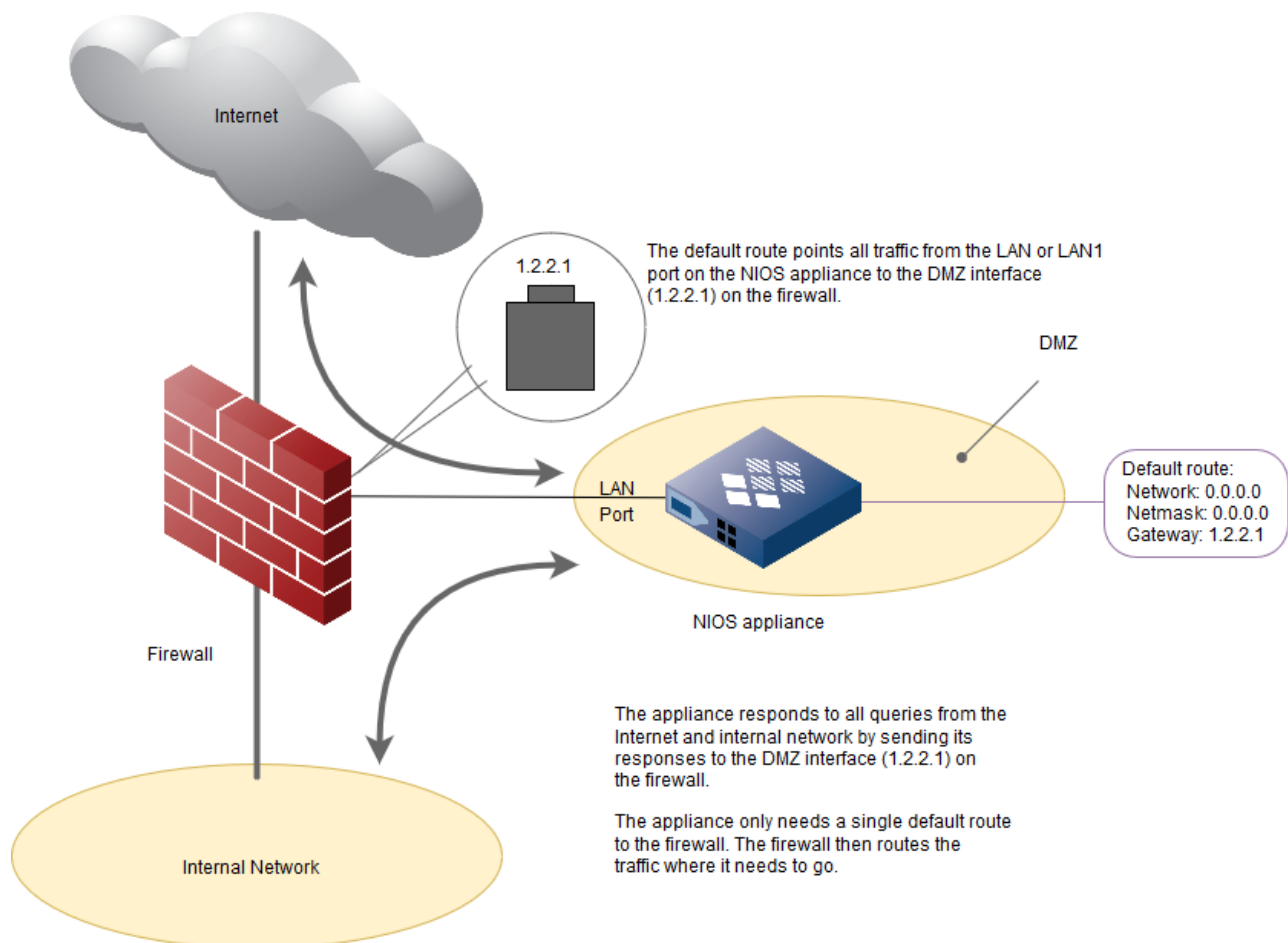
```
Copyright (c) 1999-2012 Infoblox Inc. All Rights Reserved. type 'help' for more information
```

```
Infoblox >
```

## Setting Static Routes

When you put the NIOS appliance on a segment of the network where there is a single path to and from it, a single default route is sufficient. For example, in the figure *Single Default Route* below, the appliance is in the DMZ behind a firewall and connects to the rest of the network through the DMZ interface on the firewall. For example, when hosts send DNS queries from the Internet and the internal network to the appliance and when the appliance replies to those hosts, the firewall takes care of all the routing.

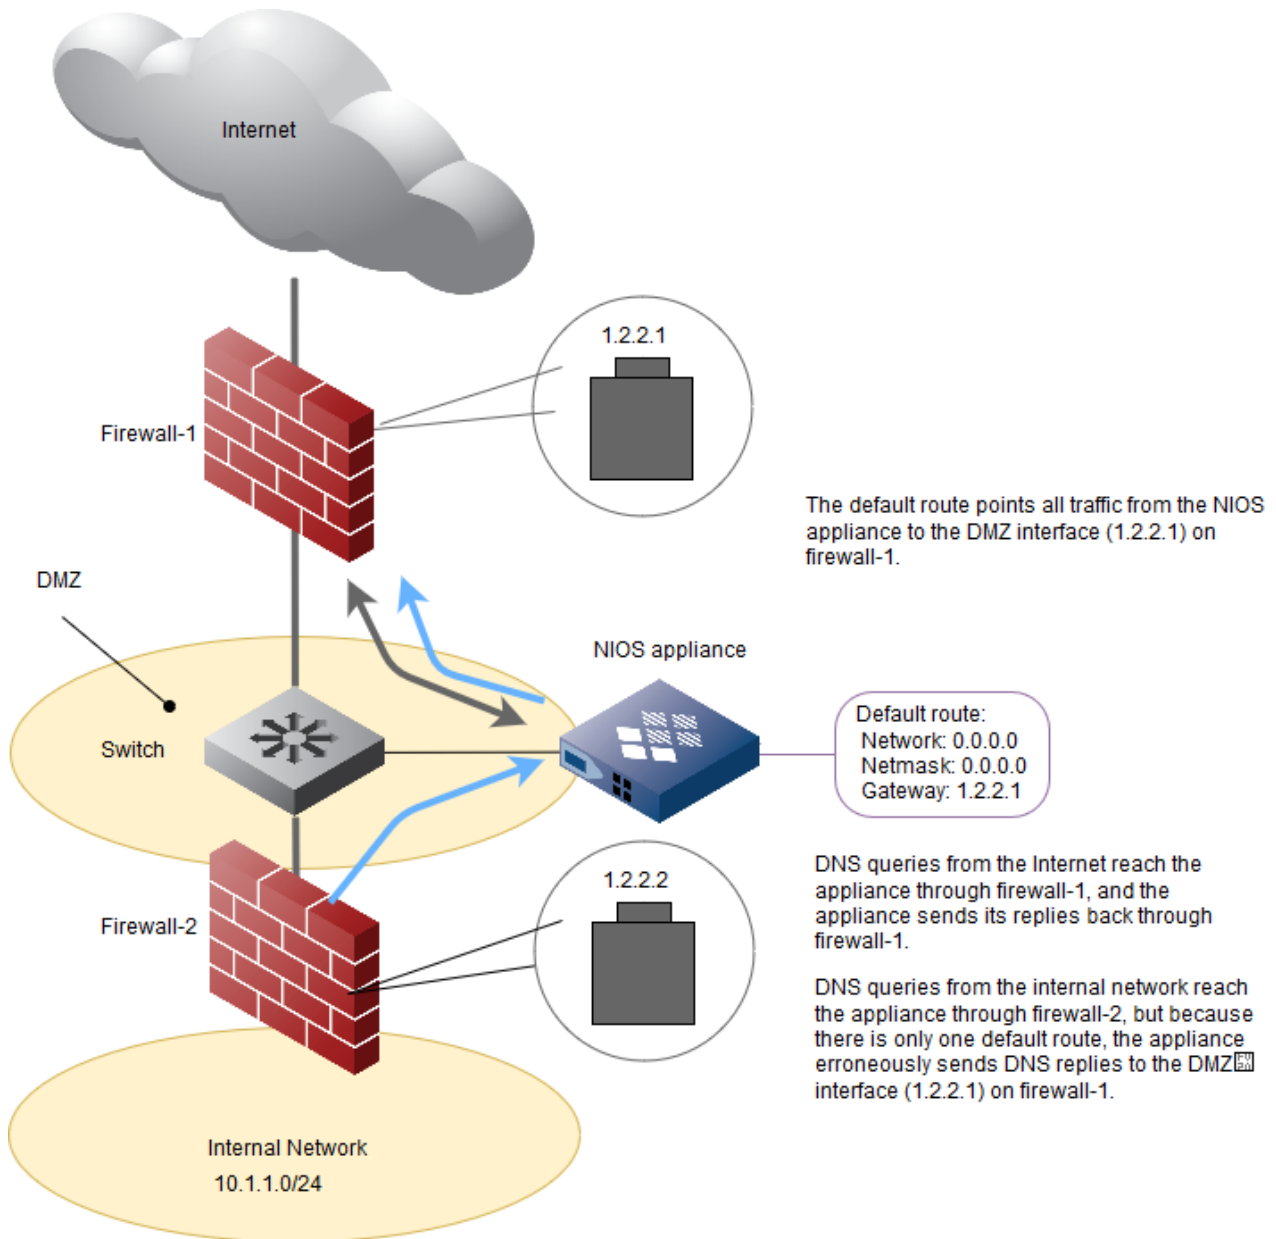
### *Single Default Route*



When the NIOS appliance is on a segment of the network where there are multiple gateways through which traffic to and

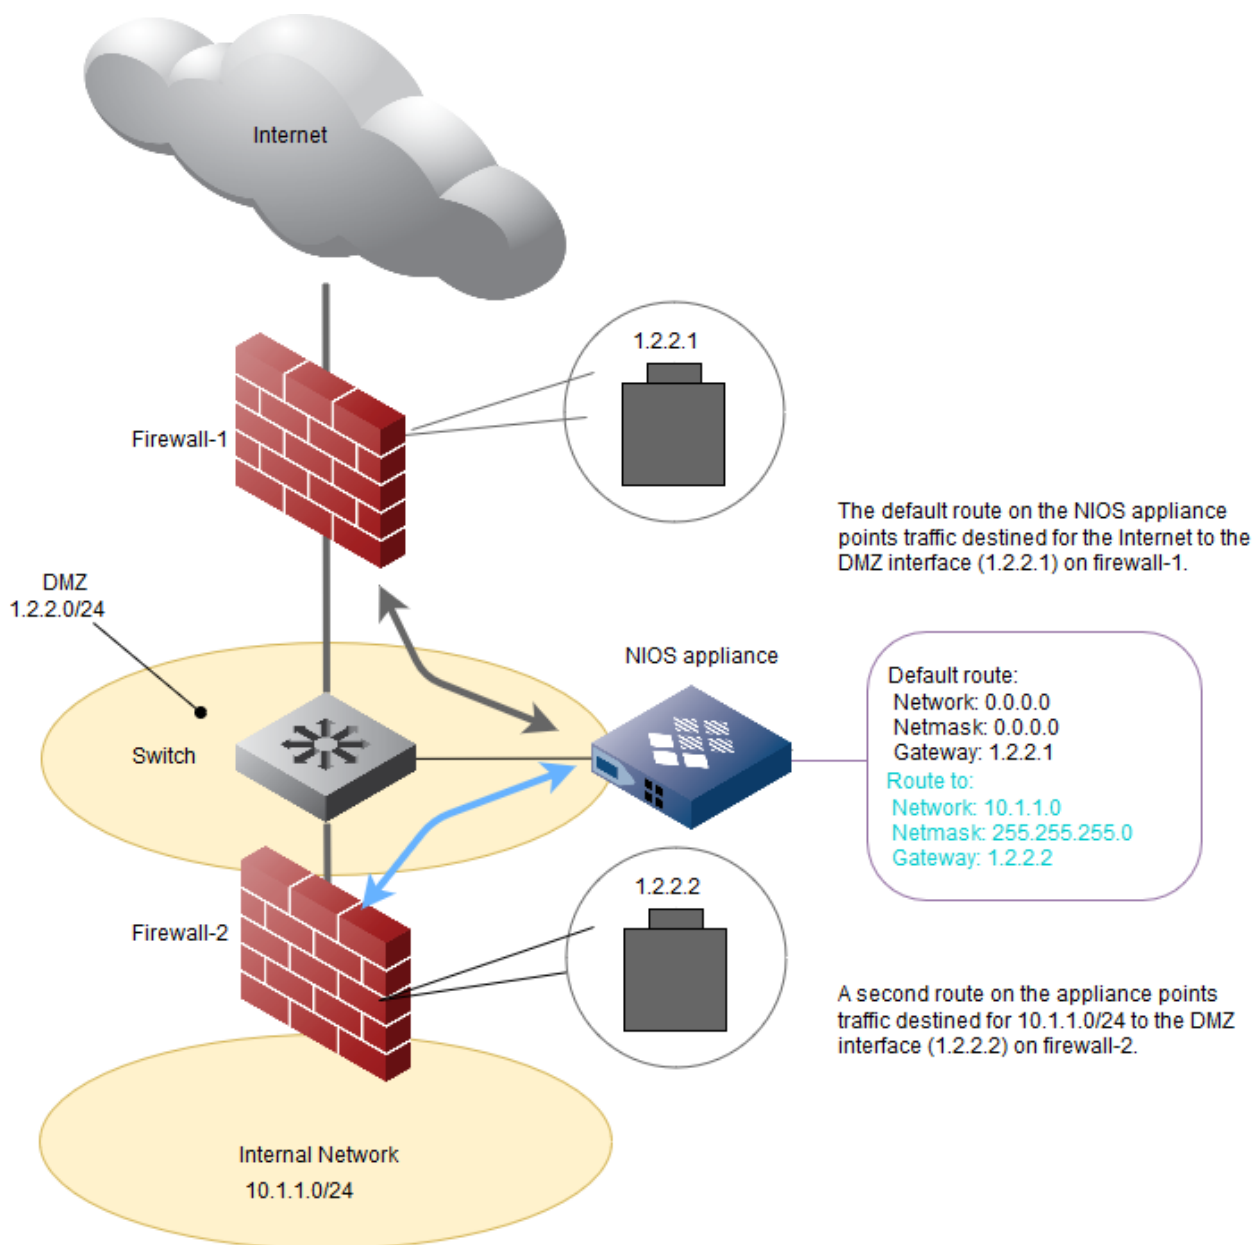
from the appliance can flow, a single default route is insufficient. For an example, see the below figure.

### *Erroneously Routed DNS Replies*



To resolve the problem illustrated in the figure *Erroneously Routed DNS Replies* above, add a second route pointing traffic destined for 10.1.1.0/24 to use the gateway with IP address 1.2.2.2 on firewall-2. This is shown in the below figure.


### *Properly Routed DNS Replies*



Whenever you want the NIOS appliance to send traffic through a gateway other than the default gateway, you need to define a separate route. Then, when the appliance performs a route lookup, it chooses the route that most completely matches the destination IP address in the packet header. When you enable the MGMT port, the gateway you reference in a static route determines which port the NIOS appliance uses when directing traffic to a specified destination.

- If a route definition references a gateway that is in the same subnet as the IP and VIP addresses of the LAN (or LAN1) and HA ports, the NIOS appliance uses the LAN (or LAN1) or HA port when directing traffic to that gateway.
- If a route definition references a gateway that is in the same subnet as the MGMT port, the NIOS appliance uses the MGMT port when directing traffic to that gateway

*Static Routes for the LAN and MGMT Ports*



Two static routes direct traffic from the NIOS appliance:

- From the LAN port (eth1, 1.2.2.5) through the gateway at 1.2.2.1 to the 10.1.1.0/24 subnet.
- From the MGMT port (eth0, 10.1.2.5) through the gateway at 10.1.2.1 to the 10.1.3.0/24 subnet.

Route Tables on the NIOS appliance

```

From LAN:
1.2.2.0/24 dev eth1 scope link
10.1.1.0/24 via 1.2.2.1 dev eth1
default via 1.2.2.1 dev eth1

From MGMT:
10.1.2.0/24 dev eth0 scope link
10.1.3.0/24 via 10.1.2.1 dev eth0
default via 10.1.2.1 dev eth0

From all:
10.1.1.0/24 via 1.2.2.1 dev eth1
10.1.3.0/24 via 10.1.2.1 dev eth0
1.2.2.0/24 dev eth1 proto kernel scope link src 1.2.2.5
10.1.2.0/24 dev eth0 proto kernel scope link src 10.1.2.5
default via 1.2.2.1 dev eth1
  
```

Note: There is a route table for each port as well as a comprehensive route table. For an HA pair, the LAN port route table is duplicated for the HA port.

In this illustration, the static routes are shown in green.

The need for routes can apply to any type of traffic that originates from the appliance, such as DNS replies, DHCP messages, SNMP traps, ICMP echo replies, Infoblox GUI management, and Grid communications. To set a static route, do the following:

- From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
- In the **Network** -> **Advanced** tab of the *Grid Member Properties* editor, click the Add icon for the **IPv4 Static Routes** table, and then enter the following:
  - Network Address:** Type the address and netmask of the remote network to which the NIOS appliance routes traffic.
  - Gateway Address:** Type the IP address of the gateway on the local subnet through which the NIOS appliance directs traffic to reach the remote network. The gateway address must meet the following requirements:
    - It must belong to a working gateway router or gateway switch.
    - It must be in the same subnet as the NIOS appliance.
 Note to consult your network administrator before specifying the gateway address for a static route on the appliance. Specifying an invalid gateway address can cause problems, such as packets being dropped or sent to an incorrect address.
- Save the configuration and click **Restart** if it appears at the top of the screen.

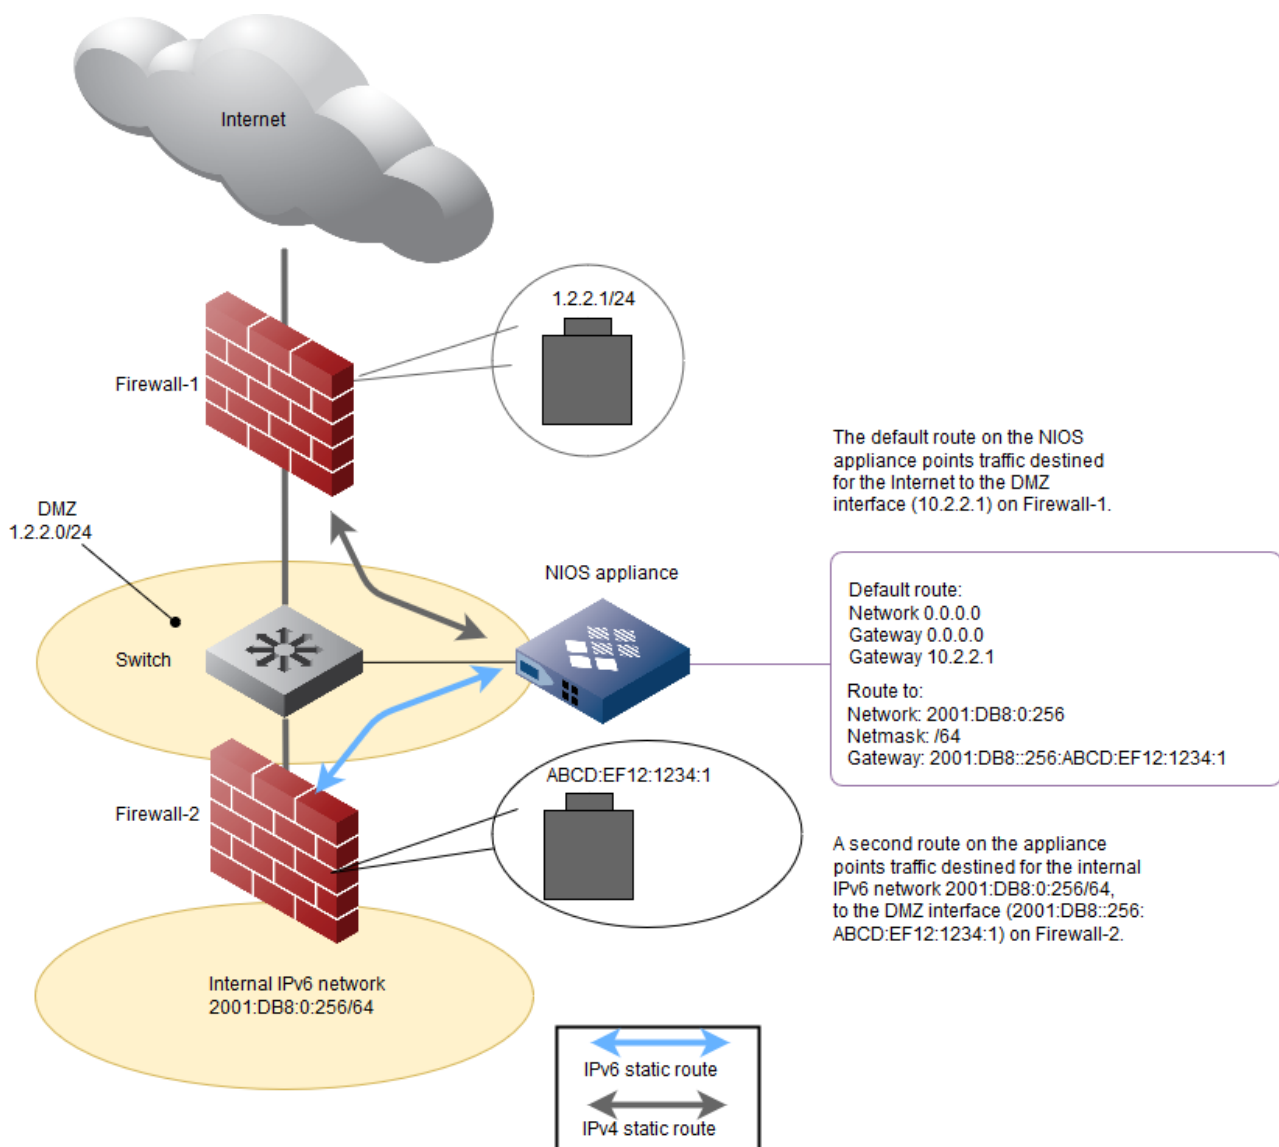
## Defining IPv6 Static Routes

Principles and applications related to IPv4 static routing in this section apply equally to IPv6. In the figure *Static Routes for IPv6 Traffic* below, a NIOS appliance supports both IPv4 and IPv6 on its LAN1 port. IPv6 is routed to the internal network while the default IPv4 route remains to the outbound 10.2.2.1 address.

You can use prefix notation to enter an IPv6 network address; the full 128-bit gateway value must be entered. To set an IPv6 static route, do the following:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. In the **Network** -> **Advanced** tab of the *Grid Member Properties* editor, click the Add icon for the **IPv6 Static Routes** table, and then enter the following:
  - **Network Address:** Type the prefix and prefix length of the remote network to which the NIOS appliance routes traffic. As an example: 2001:DB8::256/64. The double colon is required at the end of the prefix. NIOS performs validity checks on the address while it is being entered.
  - **Gateway Address:** Type the IP address of the gateway on the local subnet through which the NIOS appliance directs traffic to reach the remote network. As an example: 2001:DB8::256:ABCD:EF12:1234:1. The gateway address must meet the following requirements:
    - It must belong to a working gateway router or gateway switch.
    - It must be in the same subnet as one of the interfaces of the NIOS appliance.
    - The gateway address cannot be the same value as that for the VIP.Note to consult your network administrator before specifying the gateway address for a static route on the appliance. Specifying an invalid gateway address can cause problems, such as packets being dropped or sent to an incorrect address.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

### *Static Routes for IPv6 Traffic*



## Enabling DNS Resolution

You can specify a network server to perform domain name queries and specify up to two name servers for resolving a DNS name. You can specify the IP address of a preferred name server and that of an alternate name server, plus use a search list for performing partial name resolution.

To enable DNS resolution for a Grid or for an independent appliance or HA pair:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, expand the Toolbar and click **Grid Properties** -> **Edit**.  
**Member:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* check box, and then click the Edit icon.  
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the *Grid Properties* or *Member Properties* editor, select the **DNS Resolver** tab, and then enter the following:
  - **Enable DNS Resolver:** Select the check box to enable the NIOS appliance to send DNS queries to the preferred or alternate name servers whose IP addresses you specify in the following fields.  
 Click the Add icon and enter the IP addresses (IPv4 or IPv6) of the servers to which the appliance sends queries. The first address that you add is the primary DNS resolver and the second address that you add is the secondary DNS resolver. The appliance attempts to send queries to the servers in the order they are listed if it does not receive a response from a listed name server. To move a server up or down on the list, select it and drag it to its

new location or click the up and down arrows.

Note that if you are using the Infoblox Subscriber Parental Control feature, set the primary DNS resolver to a loopback address (127.0.0.1) to enable parental control.

- **Search List:** You can define a group of domain names that the NIOS appliance can add to partial queries that do not specify a domain name. For example, if you define a RADIUS authentication home server as "as1", and you list "corpxyz.com" and "hq.corpxyz.com" in the domain group list, then the NIOS appliance sends a query for as1.corpxyz.com and another query for as1.hq.corpxyz.com to the preferred or alternate name server. To specify domain names containing IDNs, manually convert it into punycode and specify domain names in punycode. To add a domain name, click the Add icon and type a domain name in the Search List field. To remove a domain name from the group, select it, and then click **Delete**.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Managing Licenses

You must install valid licenses for services and features to function properly in your Infoblox Grid. Different license types that are available are discussed in this topic. You can choose to obtain licenses for the desired features and services, and install them as static, dynamic, or Grid-wide licenses, depending on your network and business requirements.

After you install your licenses, you can monitor them from Grid Manager. Licenses are listed on the **Grid** tab -> **Licenses** tab. Grid Manager also displays the number of licenses that have expired and those that are expiring within the next 30 and 90 days respectively.

NIOS licenses are valid for TrinziC 2016 appliances, both physical and virtual. A NIOS license can be permanent or have an expiration date.

This topic includes the following sections:

- [License Types](#)
- [Managing Static Licenses](#)
  - [Retrieving the Hardware ID of an Appliance](#)
  - [Obtaining Static Licenses](#)
  - [Downloading License Keys](#)
- [Managing Dynamic Licenses](#)
  - [Obtaining Dynamic Licenses](#)
  - [Manually Allocating Dynamic Licenses](#)
  - [Manually Deallocating Dynamic Licenses](#)
  - [Manually Revoking Dynamic Licenses](#)
- [Managing Grid-wide Licenses](#)
  - [About the Flex Grid Activation License](#)
  - [About the Flex Grid Activation for Managed Services License](#)
  - [Obtaining Grid-wide Licenses](#)
- [Adding Permanent or Subscription Licenses](#)
- [Adding Temporary Licenses](#)
- [Viewing Licenses](#)
  - [Viewing Member Licenses](#)
  - [Viewing Dynamic Licenses](#)
  - [Viewing Grid-wide Licenses](#)
- [Backing Up Licenses](#)
- [Removing Licenses](#)
- [NIOS Licenses](#)

### License Types

Infoblox licenses are divided into the following classes:

- **Static:** Static licenses are member specific. They are installed and tied to the Grid Master or specific Grid members. Static licenses are mapped to the hardware ID for each individual NIOS or vNIOS appliance.



- **Dynamic:** These are floating licenses that are dynamically allocated to specific Grid members in the Grid. Dynamic licenses belong to a license pool, which is associated with the entire Grid. When not in use, they can be released back to the pool of licenses for further allocation. In an environment, such as a CMP (Cloud Management Platform), where you need to spin up multiple remote vNIOs appliances at different times based on business requirements, you can consider installing dynamic licenses.
- **Grid-wide:** Grid-wide licenses are associated with the entire Grid. They are not tied to any particular member. When installed, Grid-wide licenses are replicated to all members in the Grid. Although a Grid-wide license entitles all Grid members to run a particular feature, other conditions and factors determine whether the feature can be enabled on a particular member. For example, a member might not be able to run the Reporting and Analytics feature because it does not have the reporting appliance model. The currently supported Grid-wide licenses are **Security Ecosystem, Reporting Subscription, RPZ, Threat Analytics, Flex Grid Activation, and FireEye**. To configure Grid-wide licenses for RPZ, see [Grid-wide licenses for RPZ](#).

The static, dynamic, and Grid-wide licenses can have one of the following terms of duration:

- **Perpetual:** Perpetual licenses are permanent licenses that do not have expiration dates.
- **Subscription** licenses: Unlike perpetual licenses, subscription licenses come with an expiry date and will be valid till the expiry of the contract. After subscription licenses expire, the renewed licenses must be downloaded from the support portal and applied. These licenses can be static or Grid-wide licenses. For more information about Subscription Licensing, see [Licensing Requirements](#).
- **Temporary** licenses: You can enable one of several sets of temporary service licenses through the CLI command `set temp_license`. The duration for a temporary license is 60 days. They provide licensed features and functionality for the interim, while you wait for your permanent or subscription licenses to arrive. After these licenses expire, you cannot subscribe to the temporary licenses again. Note that pool licenses are not supported as temporary licenses.

Before any non-perpetual (subscription or temporary) licenses expire, an expiration warning appears during the Grid Manager login process. The warning reappears during each login until you renew the license. To renew a license, contact your accounts representative at Infoblox.



#### Note

Grid Manager does not distinguish between subscription and temporary licenses.

## Managing Static Licenses

Static licenses member-specific licenses that are mapped to the hardware IDs (also called as serial numbers) of individual NIOS or vNIOs appliances. Static licenses cannot be transferred among appliances. They come pre-installed on physical NIOS appliances according to the software packages you order at the time of purchase.

Static licenses can be perpetual or non-perpetual. A perpetual license does not have an expiration date. When you install a new static license with an expiration date beyond the existing license, it replaces the existing license. After you install static licenses on an appliance, you can view their status in Grid Manager.

### Retrieving the Hardware ID of an Appliance

After you have set up or pre-provisioned a physical appliance or vNIOs virtual appliance, obtain the hardware ID or the virtual host hardware ID in case of a vNIOs virtual appliance:

1. Log in to the NIOS CLI, or open a terminal session for the VM and open the NIOS CLI.
2. Run the `show hwid` command as shown in the following example:

```
nios-vm-ib-1420-gm > show hwid
Hardware ID: 564d41e13a1cc55affb9bad4e3b5c48a
```

Copy and paste the `Hardware ID` value for convenience.

You can also run a `show license` command to obtain the same information:

```
nios-vm-ib-1420-gm > show license
```

Version : 6.11.0-248090

Hardware ID : 564d41e13a1cc55affb9bad4e3b5c48a



#### Note

Make a note of the hardware IDs that you obtain during this procedure. Each of these unique Hardware ID values can be associated with a License Activation ID from your Contract Notification email. If a license key is installed for the current VM, that key value also appears in the output for the `show license` command.

## Obtaining Static Licenses

You can retrieve a license key for a static license either in CSV (comma separated values) format or in .txt format. It includes the following information in the format: serial number, hardware ID, license type, end date, and license string.

You can either upload the CSV file to the appliance or copy the license information from the .txt file and paste it in the **Paste License(s)** field on the **Grid** tab -> **Licenses** tab -> **Member** tab -> **Add** dialog in Grid Manager. For more information, see the [Adding Permanent or Subscription Licenses](#) section.

To obtain permanent static licenses from Infoblox:

1. Log in to the support portal at <https://support.infoblox.com> using Google Chrome for best performance.
2. Click **My Products**.  
An overview page that contains the list of appliance (hosts) maintenance contracts, software entitlements (assigned and unassigned), and subscriptions is displayed.
3. For partners only: In the **Select Account** drop-down list, select the account for which you want to obtain licenses.
4. To obtain the license for a new VM:
  - a. Click **Create Virtual Host**.
  - b. In the *Create Virtual Host* dialog box, enter the virtual host hardware ID and choose the license technology as **DDI (for NIOS and BloxOne)** or **NetMRI**.
  - c. Click **Save & Assign Software**.  
The *Manage Software of Host <host name>* page is displayed for the host you created.
  - d. In the *Software available to be assigned to Host* section, enter the quantity to be assigned in the **Qty to Assign** field of the applicable license.
  - e. Click **Assign Software** to assign the licenses.

SKU	SKU DESCRIPTION	MAINT. END DATE	ATTRIBUTE	ACTIVATION	TOTAL QTY	AVAILABLE QTY	QTY TO ASS...
8-SWTL-ADNS	Tritonic Software Module Subscription, Advanced DNS Protection	Nov 12, 2020		0346809	2	2	-- 1 +
1G-SW-PT	Advanced Software Module Perpetual License, ADNS	Oct 9, 2018		000572280003003-1	1	1	-- 1 +

- f. Click the **Download License Keys** button to download the licenses.
5. To obtain the license for a new physical appliance:
    - a. Click the **Hosts** tab.
    - b. In the search box, enter criteria to search for your appliance.
    - c. In the **Manage** column of your appliance, select **Manage Software** from the drop-down list.  
The *Manage Software of Host <host name>* page is displayed.
    - d. Assign licenses to the appliance or unassign them if not needed:
      - i. In the *Software available to be assigned to Host* section, enter the quantity to be assigned in the **Qty to Assign** field of the applicable license, and click **Assign Software**.
      - ii. In the *Software currently assigned to Host* section, enter the quantity to be unassigned in the **Qty to Unassign** field of the applicable license, and click **Unassign Software**.
    - e. Click the **Download License Keys** button to download the licenses.
  6. To obtain the license for an existing physical or a virtual appliance for which you are unsure of license entitlements or want to obtain a new listing of your license keys:
    - a. Click the **Hosts** tab.
    - b. In the search box, enter criteria to search for your appliance.
    - c. In the **Manage** column of your appliance, select **Download License Keys** from the drop-down list.

- d. Additionally, you may assign or unassign licenses and download them from the *Manage Software for Host <host name>* page as explained in step 5.



#### Note

Each VM License Activation ID should have a Hardware ID associated with it. As you install and spin up each virtual machine, establish written records for each Hardware ID with the VM License Activation IDs in a one-to-one ratio. These value pairs are necessary should you need to contact Infoblox Technical Support.

## Downloading License Keys

You can download the license information either in CSV or TXT format. The following options are available in the [Support](#) portal:

- To download the license information for an individual appliance, click **Download License Keys** that is available on the **My Products** tab -> **Hosts** tab -> drop-down list, or on the *Manage Software for Host <host name>* page.
- To download the DDI license information for all your appliances at a time, click **Download All DDI License Keys** that is available on the **My Products** tab -> **Hosts** tab.

In the *Download License Keys* dialog that opens, complete the following steps:

1. Select between **TXT** or **CSV** format.
2. Click **Download**.

**Download License Keys**  
These file types can be imported or pasted into the Infoblox appliance.

Download File Format

TXT  CSV

Serial #	Activation Id	Quantity	Maint. End Date	Feature	License Key
	0369508	1		NIOS	
	0369509	1		DHCP	
	0369509	1		DNS	

← Prev Page 1 of 1 - Total Records: 3 Next →

Cancel Download

3. In the *Save as* dialog, choose the location where you want to save the file.
4. Click **Save**.

## Managing Dynamic Licenses

Dynamic licenses are multiple licenses you obtain for NIOS appliances, both physical and virtual. You can install dynamic licenses in advance for different services and features and deploy them on demand, based on your business needs. The appliance stores these licenses in their respective license pools and allocates them when you deploy vNIOS virtual appliances. When you remove a vNIOS appliance from your Grid, the applicable licenses are automatically released back to the license pool. The appliance adjusts the total number of licenses accordingly. You can use dynamic licenses to automatically pre-provision and deploy vNIOS appliances. For more information about this feature, see [About Elastic Scaling](#).

For a dynamic license, you can either upload the license file to the appliance or copy the information and paste it in the **Paste License(s)** field on the **Grid** tab -> **Licenses** tab -> **Pool** tab in Grid Manager. You must copy the entire string and save it to the text field. When you restore or perform a factory reset operation, you will lose the existing dynamic licenses.

License pools for a particular license type are an aggregate of the subpools that contain the respective license type. Subpools are either perpetual or non-perpetual with different expiration dates. NIOS displays non-perpetual and perpetual subpools with different expiration dates within the same pool.

Dynamic licenses can be allocated from a pool based on the aggregate install count. You can allocate licenses from a pool as long as the allocated count of perpetual and non-perpetual licenses does not exceed the count of perpetual and active non-perpetual licenses that are installed.



#### Note

You must install both the Grid and vNIOS licenses on a vNIOS appliance for it to join the Grid. You can add other licenses such as DNS, DHCP, or Cloud Platform depending on how you deploy your vNIOS virtual appliances.

## Obtaining Dynamic Licenses

When you purchase licenses for specific features and services, you may deploy the licenses as dynamic licenses. License information is stored in a license file with the following information for each license: LSN-P (License Serial Number - Proxy), LK-P (License Key Proxy, LSN (License Serial Number) and LK (License Key). License keys are generated based on the LPC\_UID (License Pool Container Unique ID) of the Grid Master. You must first obtain the LPC\_UID of the Grid, also called the Grid license UID from the Grid Master and then contact Infoblox Technical Support to obtain dynamic licenses.

Infoblox stores information related to dynamic licenses in a license file. When you install dynamic licenses, you must upload the entire license file to the Grid Master, as described in the [Adding Permanent or Subscription Licenses](#) section. To obtain dynamic licenses:


1. Log in to Grid Manager and obtain the LPC\_UID from the Grid Master:
  - a. From the **Grid** tab, select the **Licenses** tab -> **Pool** tab, and then click **Export All Licenses** in the Toolbar.
  - b. Save the CSV file.
  - c. Open the CSV file. The LPC\_UID is displayed in the **SIGNATURE** row. Copy this ID. You will need this ID when contacting Infoblox Technical Support.

Note that you can also use the show `license_pool_container` CLI command to display the `LPC_UID`.

2. Contact Infoblox Technical Support or your Infoblox representative to obtain the dynamic licenses.

## Manually Allocating Dynamic Licenses

To allocate a particular feature license to a NIOS member:

1. From the **Grid** tab, select the **Licenses** tab -> **Pool** tab.
2. Select the feature license and click the Action icon  next to the license, and then select **Allocate** from the list.
3. In the *Allocate Pool Licenses* dialog, click **Select Member Node**.
4. From the *Member Selector*, click the FQDN of the NIOS member on which you want to install the feature license. The appliance displays the member in the **To this member:** field.
5. Click **Allocate** to install this license on the selected member or click **Cancel** to cancel this action. NIOS allocates the license and adjusts the numbers in the **Assigned** and **Available** fields on the **Pool** tab.

## Manually Deallocating Dynamic Licenses

To deallocate a particular feature license from a NIOS member:

1. From the **Grid** tab, select the **Licenses** tab -> **Member** tab.
2. Select the license and click the Action icon next to the checkbox, and then select **Deallocate** from the list.
3. In the *Deallocate Pool Licenses* dialog, select the licenses you want to deallocate from the Pool Licenses table.
4. Click **Deallocate** to remove the licenses from the members or click **Cancel** to cancel this action. NIOS deallocates the license from the respective members and adjusts the numbers in the **Assigned** and **Available** fields on the **Pool** tab.

## Manually Revoking Dynamic Licenses

To remove a particular feature license from the pool:

1. From the **Grid** tab, select the **Licenses** tab -> **Pool** tab.
2. Select the feature license and click the Action icon next to the checkbox, and then select **Revoke** from the list.
3. In the *Remove Licenses* dialog, click the **Number** field and enter the number of licenses you want to remove for the selected license type.
4. Click **Remove Licenses** to remove the licenses or click **Cancel** to cancel this action.  
NIOS removes the licenses and adjusts the numbers in the **Assigned** and **Available** fields on the **Pool** tab.

## Managing Grid-wide Licenses

NIOS supports Grid-wide licenses that are valid across the entire Grid. You can obtain Grid-wide licenses for features listed in this section. Although a Grid-wide license entitles all Grid members to run a particular feature, other conditions and factors determine whether the feature can be enabled on a particular member. For example, a member might not be able to run the Reporting and Analytics feature because it does not have the reporting appliance model.

For a Grid-wide license, you can either upload the license file to the appliance or copy the information and paste it in the text field on the **Paste License(s)** field on the **Grid** tab -> **Licenses** tab -> **Grid Wide** tab of Grid Manager. You must copy the entire string (license type, expiry date, and license string). The **Flex Grid Activation** license is bundled with the following licenses: Grid, Unbound, DNS Cache Acceleration, DNS, DHCP, DNS Traffic Control, Response Policy Zone, Software Threat Protection, Threat Protection Update, DNSFW, NXDOMAIN Redirect, Dual Engine DNS (only for recursive DNS), Microsoft Management, Threat Analytics, Security Ecosystem, and Cloud Network Automation.

Note the following about Grid-wide licenses:

- When you restore or perform a factory reset operation, you will lose the existing Grid-wide licenses.
- NIOS restores any Grid-wide licenses that are present in a restored database.
- When you restore a database from another Grid, NIOS replaces the UID of the Grid and the Grid-wide licenses from the other Grid.
- Static and Grid-wide licenses of the same type can co-exist in the same Grid.
- If a member is pre-provisioned for a specific feature, it is allowed to join the Grid that has the Grid-wide license for that feature even if the member does not have a license for that feature.
- To install a Cloud Network Automation license on an IB-FLEX appliance, it must be set up as a Grid Master.

## About the Flex Grid Activation License

**Flex Grid Activation** is a permanent Infoblox license that you can implement as a NIOS Grid-wide license to enable the following features at the same time:

- Grid
- Unbound
- DNS Cache Acceleration
- DNS
- DHCP
- DNS Traffic Control
- Response Policy Zone
- Software ADP
- Threat Protection Update
- DNSFW
- NXDOMAIN Redirect
- Query rewrite
- Threat Analytics
- Security Ecosystem
- Captive Portal
- Microsoft Management
- Cloud Network Automation

For more information, see [About IB-Flex](#). After you install the **Flex Grid Activation** license, you can view it on the **Grid** tab -> **Licenses** tab -> **Grid Wide** tab of Grid Manager.

You can install the **Flex Grid Activation** license on a Grid Master even if IB-FLEX is not a part of the Grid, but this license is effective only for an IB-Flex member. The Infoblox License Portal allows you to acquire any number of **Flex Grid Activation** license keys for each individual Grid. To install a temporary license, use the [set temp\\_license](#) CLI command.

## About the Flex Grid Activation for Managed Services License

Flex Grid Activation for Managed Services is a combination of all existing Infoblox licenses that you can implement as a NIOS Grid-wide license. If you install this license, you do not need to install each license separately. The Flex Grid Activation for Managed Services license comprises the following licenses:

- Grid
- Unbound
- DNS Cache Acceleration
- DNS
- DHCP
- DNS Traffic Control
- Response Policy Zone
- NXDOMAIN Redirect
- Dual Engine DNS (only for recursive DNS)
- Software Threat Protection
- Threat Protection Update
- Threat Analytics
- Security Ecosystem
- Microsoft Management
- Cloud Network Automation

You can install the Flex Grid Activation for Managed Services license only on a Grid Master or a standalone appliance. To install a temporary license, use the [set temp\\_license](#) CLI command.



### Note

If you install the Flex Grid Activation for Managed Services license, you cannot install the Flex Grid Activation license and if you install the Flex Grid Activation license, you cannot install the Flex Grid Activation for Managed Services license.

## Obtaining Grid-wide Licenses

When you purchase licenses for specific features and services, the licenses may be Grid-wide licenses. License information is stored in a license file with the following information for each license: LSN-P (License Serial Number - Proxy), LK-P (License Key Proxy, LSN (License Serial Number) and LK (License Key).

License keys are generated based on the license UID of the Grid Master. You must first obtain the UID of the Grid and then contact Infoblox Technical Support to obtain the Grid-wide licenses.



### Note

Ensure that you obtain the license UID of the Grid Master. If you use the license UID of a Grid member or an appliance that has not yet joined the Grid, you might not be able to properly install the Grid-wide license.

The license file (CSV) format for Grid-wide licenses are as follows:

```
GRID_WIDE,license_uid,license_type,[expiry_date],license_string
```

Infoblox stores information related to Grid-wide licenses in a license file. When you install Grid-wide licenses, you must upload the entire license file to the Grid Master, as described in the *Adding Permanent or Subscription Licenses* section.

To obtain Grid-wide licenses:

1. Log in to Grid Manager and obtain the license UID of the Grid Master as follows:
  - a. From the **Grid** tab, select the **Licenses** tab -> **Grid Wide** tab, and then click **Export All Licenses** from the Toolbar.
  - b. Save the CSV file.
  - c. Open the CSV file. The UID is displayed in the **SIGNATURE** row. Copy this ID. You will need this ID when contacting Infoblox Technical Support.

Note to obtain the UID of the Grid, you can use the `show license_uid` command on the Grid Master. For more information, refer to the *Infoblox CLI Guide*.

2. Contact Infoblox Technical Support or your Infoblox representatives to obtain the Grid-wide licenses.

## Adding Permanent or Subscription Licenses

To add permanent or subscription licenses:

1. From the **Grid** tab, select the **Licenses** tab -> **Member** tab or **Pool** tab or **Grid Wide** tab and click the Add icon.
2. Select one of the following:
  - **Upload License File:** Use this method to upload static, dynamic, and Grid-wide license files. Click **Select File**, navigate to the license file, and click **Open**.
  - **Paste License(s):** Paste the license key in this text field. You must paste the entire string in CSV format: serial number, hardware ID, license type, end date (in the mm/dd/yyyy format), and license string. If you are pasting multiple licenses, start each string on a new line. For example, 564daaa3ef07e648a563434a2b412834,564daaa3ef07e648a563434a2b412834,DNS,12/05/2018,EwAAADU+PZEOjFsMAXxw4nlfHyZ6tdw
3. Click **Save License(s)**.  
The appliance validates the license and adds it to the table. Close the browser window and log in to Grid Manager. If you are activating licenses for an HA pair, you must follow this procedure for both nodes.



### Note

- Once the NIOS license is installed, you must wait for NIOS to restart (if prompted) before you install other licenses.
- To transfer licenses between vNIOS on VMware appliances, refer to the *Infoblox Installation Guide for vNIOS for VMware*.

## Adding Temporary Licenses

Use the `set temp_license` command to generate and install temporary licenses. This can provide licensed features and functionality for the interim, while you wait for your permanent license to arrive.

To generate a temporary license:

1. Log in to the NIOS appliance through a remote console window.  
For more information on how to open a remote console window, refer to the *Infoblox CLI Guide*.
2. At the Infoblox command prompt, enter `set temp_license`.  
The appliance lists the available licenses, and you select those you need. For more information, see [set temp\\_license](#).
3. Enter the number of licenses you want to install.
4. Confirm the selection when prompted, and the following message appears:  
`Temporary license is installed.`



### Note

Once the NIOS license is installed, you must wait for NIOS to restart before you install other licenses.

## Viewing Licenses

If the appliance is part of a Grid, you must log in to the Grid Master to view license information from Grid Manager. If the appliance is an independent appliance, log in to System Manager on the appliance. If you have transferred licenses from one vNIOS appliance to another, you can view information about the new and replaced licenses.

Grid Manager displays licenses on the **Grid** tab -> **Licenses** tab. You can view license information for all static and dynamic licenses (including temporary licenses) on the **Member** tab, and a summary of dynamic licenses on the **Pool** tab.



## Viewing Member Licenses

To view the information about active licenses, including static and dynamic licenses that are currently assigned or allocated to NIOS and vNIOS appliances:

1. Log in to Grid Manager on the Grid Master or System Manager on an independent appliance.
2. Select the **Grid** or **System** tab -> **Licenses** tab -> **Member** tab. The appliance displays the following information:
  - **Type of License:** The license category. This can be **Static**, **Dynamic**, **Grid Wide**, or **Paid NIOS**. Static licenses are individual licenses you obtain and are currently assigned to specific appliances. These licenses are tied to specific hardware IDs and you cannot deallocate them. Dynamic licenses are pooled licenses that support the Elastic Scaling feature, which enables central tracking, granting, and management of NIOS feature licenses for vNIOS entities in the Grid. You can manually allocate and deallocate dynamic licenses. When installed, Grid-wide licenses are replicated to all members in the Grid. The currently supported Grid-wide licenses are **Security Ecosystem**, **Reporting Subscription**, **RPZ**, **Flex Grid Activation** and **FireEye**. Paid NIOS represents the pay-as-you-go licensing model for vNIOS virtual appliances. In the AWS Marketplace, when you use the Paid NIOS model to launch the vNIOS for AWS virtual appliance, the virtual appliance comes pre-installed with the following permanent licenses: **vNIOS**, **Grid**, **DNS**, and **CNA** (Cloud Network Automation). As long as the virtual appliance is up and running, you can use the NIOS features that these licenses provide. You cannot add, delete, import, or export Paid NIOS licenses. For information about the Paid NIOS in AWS Marketplace, refer to the *Installation Guide for vNIOS for AWS*.
  - **Feature:** Indicates the features for which the license was installed. For example, if the license was installed for DNS service, this shows **DNS**.
  - **Name:** The FQDN of the appliance.
  - **HA:** Indicates whether the appliance is an HA pair.
  - **IPv4 Address:** The IPv4 address of the appliance, if applicable.
  - **IPv6 Address:** The IPv6 address of the appliance, if applicable.
  - **Hardware ID:** The unique hardware ID of the appliance. The ID is highlighted in red if the license on the appliance was removed.
  - **Serial Number:** The serial number of the appliance.
  - **Type Context:** Depending on the license type, this field displays the attribute (such as **Model**) that the license controls. This field is blank if the license does not control any attribute type. This field can display one of the following:
    - **Leases:** Indicates that this DHCP license supports a specific number of DHCP leases. The number of leases supported is displayed in the **Type Details** field.
    - **Model:** Indicates that this vNIOS license supports a specific vNIOS virtual appliance model. The model supported is displayed in the **Type Details** field.
  - **Type Details:** Information about the attribute type that the license monitors. This field can display the following information for each attribute:
    - **Leases:** The number of DHCP leases that the DHCP license supports.
    - **Model:** The model of the NIOS virtual appliance, such as IB-V1410 or IB-V2215.
  - **Expiration:** The expiration date of the license.
  - **Replaced Hardware ID:** The hardware ID of the appliance whose license was removed.

## Viewing Dynamic Licenses

You can install dynamic licenses on the Grid Master for future vNIOS appliance deployments. When you install dynamic licenses, the Grid Master store them in a license pool. You can view these licenses and evaluate license usage for vNIOS virtual appliances.

For a particular feature, such as DNS or DHCP, Grid Manager displays dynamic licenses in sub pools so you can view the number of permanent licenses, the number of licenses that are expiring on a particular date, and those that are expired. NIOS highlights the licenses that are going to expire soon using a yellow background. Licenses that are approaching expiry are highlighted with a pink background.



### Note

The overall license status for a particular feature reflects the most critical status among all licenses in the pool. For example, if there are expired licenses in the pool, the overall status for this license type appears as expired.



To view dynamic licenses in the pool and their usage information:

1. Log in to Grid Manager on the Grid Master.
2. Select the **Grid** tab -> **Licenses** tab -> **Pool** tab. The appliance displays the following information:
  - **Feature:** The feature for which you have obtained the license.
  - **Installed:** The number of licenses that have been installed for the specified feature.
  - **Assigned:** The number of licenses currently allocated to vNIO virtual appliances.
  - **Available:** The number of licenses that are currently available for the specified feature.
  - **License Model:** For vNIO license only. Indicates the model type of the vNIO virtual appliance. Note that the vNIO license you install on the vNIO appliance must match the appliance model. You can use vNIO license that has a higher capacity on a vNIO appliance that has a smaller capacity, but not vice versa. For example, you can install a CP-V1405 license on a CP-V805 or CP-V1405 appliance, but not on a CP-V2205.
  - **Limit Context:** Depending on the license type, this field displays the attribute (such as **Model**) that the license controls. This field is blank if the license does not control any attribute type. This field can display one of the following:
    - **Leases:** Indicates that this DHCP license supports a specific number of DHCP leases. The number of leases supported is displayed in the **Type Details** field.
    - **Model:** Indicates that this vNIO license supports a specific vNIO virtual appliance model. The model supported is displayed in the **Type Details** field.
  - **Limit Value:** Information about the attribute type that the license monitors. This field can display the following information for each attribute:
    - **Leases:** The number of DHCP leases that the DHCP license supports.
    - **Model:** The model of the vNIO virtual appliance, such as IB-V410 or IB-V2215.
  - **Expiration:** The expiration date and time of the license. It displays **Permanent** for permanent licenses and **Expired** for licenses that expired.
3. Click the arrow mark next to the checkbox of a specific feature to view the list of licenses and their respective expiration dates.

## Viewing Grid-wide Licenses

Grid Manager displays the licenses that are configured for the respective Grid. To view Grid-wide licenses:

1. Log in to Grid Manager on the Grid Master.
2. Select the **Grid** tab -> **Licenses** tab -> **Grid Wide** tab. The appliance displays the following information:
  - **Feature:** The feature for which you have obtained the license.
  - **Limit Context:** Depending on the license type, this field displays the attribute (such as **Model**) that the license controls. This field is blank if the license does not control any attribute type. This field can display one of the following:
    - **Leases:** Indicates that this DHCP license supports a specific number of DHCP leases. The number of leases supported is displayed in the **Type Details** field.
    - **Model:** Indicates that this vNIO license supports a specific vNIO virtual appliance model. The model supported is displayed in the **Type Details** field.
  - **Limit Value:** Information about the attribute type that the license monitors. This field can display the following information for each attribute:
    - **Leases:** The number of DHCP leases that the DHCP license supports.
    - **Model:** The model of the vNIO virtual appliance, such as IB-V1410 or IB-V2215.
  - **Expiration:** The expiration date and time of the license. It displays **Permanent** for permanent licenses.

You can do the following:

- Delete a Grid-wide license as explained in the *Removing Licenses* section.
- Click the Export icon to export the list of licenses to a .csv file.
- Click the Print icon to print the list of licenses.

To search for specific licenses on the **Member**, **Pool** and **Grid Wide** tabs, you can use filters and the **Go to** function (the **Go to** function is not supported in the Pool tab) to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches. You can also create a quick filter to save frequently used filter criteria. For information about Using Quick Filters, see [Using Quick Filters](#).

## Backing Up Licenses

You can back up all static licenses, dynamic licenses added in the license pool container, and Grid-wide licenses in case you need to re-install them at a later time. Infoblox recommends backing up the licenses before removing any of them.



### Note

Dynamic licenses are not exported to this file. Dynamic licenses are automatically released and returned to the license pool on the Grid Master when a member leaves the Grid. Unallocated dynamic Licenses are available for allocations.

When you back up the licenses, Grid Manager creates a CSV file that lists the following information for each license: serial number, hardware ID, license type, end date, and license string.

To back up licenses:

1. From the **Grid** tab, select the **Licenses** tab.
2. Click **Export All Licenses** in the toolbar. Grid Manager generates a CSV file that contains all the licenses. Depending on the browser you use, you can then open the file or save it to a specified location.

## Removing Licenses

You can remove licenses and reset a NIOS appliance to its factory default settings. For example, if you have a NIOS appliance running the DNSone package with the Grid upgrade, but you want to use it as an independent appliance, you can remove the Grid license. Infoblox recommends that you back up licenses before removing them, in case you decide to re-install them at a future time.

When you remove a Grid-wide license, the licensed feature is deactivated on all the members that are affected by the license. On the other hand, when you remove a Grid member from the Grid, the member no longer has the ability to run the feature associated with the Grid-wide license now that is it not part of the Grid.



### Note

- Exercise caution when removing licenses; you may render an appliance unusable by removing the wrong license. Other feature sets may be affected once you remove a license; for example, removing licensing for DNS and DHCP services will also disable task packs on the **NIOS Dashboards – > Tasks** page.
- Infoblox does not recommend deleting CNA temporary licenses as it may result in unexpected behavior. For example, after deleting a CNA temporary license, running a vRA workflow to create a reservation on a vNIOS setup followed by a vDiscovery job, can result in a VM affiliation conflict. To clear the conflict, as a workaround, you must install a CNA permanent license, and then delete it. For assistance, contact the Infoblox Technical Support.

To remove an active license:

1. From the **Grid** tab, select the **Licenses** tab -> **Member** tab or **Grid Wide** tab.
2. Select the license and click the Delete icon.  
Check the license that you are about to remove. Note that removing the wrong license can render an appliance unusable.
3. Click **Yes** when the confirmation dialog appears.
4. Close the browser window and log in to the Infoblox GUI.

## NIOS Licenses

The following table lists NIOS licenses and the behavior when the subscription licenses expire:

License	Supported Features	When the License Expires
Advanced DNS Protection (ADP) / Threat Protection	Advanced DNS Protection / Threat protection	Existing functionality continues to work as is. You can add new custom rules and publish. You may not be able to upload new ruleset, but old rulesets remain functional.
Software Advanced DNS Protection / Threat Protection (software add-on)	Virtual advanced DNS protection / Threat protection	
Threat Protection Update	Receive and update threat protection rules and rule updates	
Threat Protection and Threat Protection Update	<ul style="list-style-type: none"> <li>Threat protection</li> <li>Elastic scaling</li> </ul>	
Security Ecosystem	Outbound Notification	Existing functionality continues to work as is with existing and new endpoints and notifications. Integrations such as the FireEye feed will stop functioning as it is dependent on RPZ.
FireEye	FireEye	Existing functionality continues to work as is.
DNS Cache Acceleration (DCA)	DNS Cache Acceleration	Existing functionality continues to work as is. Changes and new additions work fine.
DNS	DNS	Existing functionality continues to work as is. Allows addition of new zones and networks.
DNS Traffic Control (DTC)	DNS Traffic Control	Existing functionality continues to work as is.
Threat Analytics	Threat Insight/Threat Analytics	Existing functionality continues to work as is. New RPZ may not take effect as Named server does not restart after license expiry.
DHCP	DHCP	Existing functionality continues to work as is. Allows addition of new zones and networks.
Microsoft Management (MS MGMNT)	<ul style="list-style-type: none"> <li>Management of Microsoft DNS and DHCP servers from Infoblox Grid Manager.</li> <li>Synchronization of DNS and DHCP data to the Grid database.</li> </ul>	New data (new zone) added on Microsoft server is not synchronized to NIOS.
*Cloud Network Automation	Cloud network automation	Existing functionality continues to work as is.

License	Supported Features	When the License Expires
Cloud Platform	Cloud Platform	Existing functionality continues to work as is. The Cloud API is available. In Grid Manager, you can manage cloud objects from the <b>Grid</b> -> <b>Grid Manager</b> -> <b>Cloud-API</b> tab. You may not be able to add an A record on a zone that has Cloud Platform members assigned to it.
Response Policy Zone (RPZ)	Response Policy Zone	RPZ Feed zones and zone transfer continue to work. RPZ feeds will stop.
Dual Engine DNS (only for recursive DNS)	Unbound DNS	Functionality continues to work as is.
NIOS/vNIOS	NIOS services	Functionality continues to work as is.
Grid	<ul style="list-style-type: none"> <li>When installed on a member that has no other Infoblox licenses installed, you can Join the member to a Grid.</li> <li>When installed on a Grid Master, you can get the UID required to obtain dynamic licenses.</li> </ul>	Functionality continues to work as is with existing members. You may not be able to add new members to the Grid.
Query Redirection	Query Redirection/ NXDOMAIN Redirect	Functionality continues to work as is.
Reporting	Reporting	Reporting functionality continues to work as is and reports update. If the license related to FireEye expires, old data does not show in the <i>FireEye Alerts</i> report.
Multi-Grid Management (MGM)	Multi-Grid Management	Functionality continues to work as is.
Network Discovery (ND)	Network Discovery	Functionality continues to work as is.

\* To install a Cloud Network Automation license on an IB-FLEX appliance, it must be set up as a Grid Master.  
Licenses that are bundled with multiple licenses:

License	Supported Features	When the License Expires
DNSone includes the following licenses: <ul style="list-style-type: none"> <li>DNS</li> <li>DHCP</li> </ul>	<ul style="list-style-type: none"> <li>DNS</li> <li>DHCP</li> </ul>	-

License	Supported Features	When the License Expires
<p>The Flex Grid Activation license includes the following licenses:</p> <ul style="list-style-type: none"> <li>• Grid</li> <li>• Unbound</li> <li>• DNS Cache Acceleration</li> <li>• DNS</li> <li>• DHCP</li> <li>• DNS Traffic Control</li> <li>• Response Policy Zone</li> <li>• NXDOMAIN Redirect</li> <li>• Dual Engine DNS (only for recursive DNS)</li> <li>• Software Threat Protection</li> <li>• Threat Protection Update</li> <li>• Threat Analytics</li> <li>• Security Ecosystem</li> <li>• Microsoft Management</li> <li>• *Cloud Network Automation</li> </ul>	<ul style="list-style-type: none"> <li>• Grid</li> <li>• Unbound</li> <li>• DNS Cache Acceleration</li> <li>• DNS</li> <li>• DHCP</li> <li>• DNS Traffic Control</li> <li>• Response Policy Zone</li> <li>• Software ADP</li> <li>• Threat Protection Update</li> <li>• DNSFW</li> <li>• NXDOMAIN Redirect</li> <li>• Query rewrite</li> <li>• Threat Analytics</li> <li>• Security Ecosystem</li> <li>• Captive Portal</li> <li>• Microsoft Management</li> <li>• Cloud Network Automation (only if IB-FLEX is used as a Grid Master).</li> </ul>	<p>Functionality continues to work as is for all features except for the Grid license, where functionality works for existing members. You may not be able to add new members to the Grid.</p>
<p>The Flex Grid Activation for Managed Services license installed only on Grid Master or a standalone appliance, includes the same set of licenses bundled with the Flex Grid Activation license.</p>	<p>Same features as the Flex Grid Activation license.</p>	

\* To install a Cloud Network Automation license on an IB-FLEX appliance, it must be set up as a Grid Master.



**Note**

NTLMv2 is the only authentication method supported for Microsoft servers managed from Infoblox Grids. For information about managing Microsoft Windows servers from Grid Manager, see [Managing Microsoft Windows Servers](#).

## About IB-FLEX

IB-FLEX is a virtual platform that is scalable based on the resource that you allocate to the virtual machine. NIOS automatically detects the capacity of the virtual machine and scales it to the appropriate platform after you provision the IB-FLEX member.

You must first install the Grid license on a non IB-FLEX appliance that is designated as the Grid Master to allow members to join the Grid, even if you have already installed an **Flex Grid Activation** license. This license does not affect a non IB-FLEX Grid Master.

An IB-FLEX appliance designated as a member does not require any license, either Grid or vNIOS, while joining the Grid. When you register an IB-FLEX member, the appliance checks for the Grid (enterprise) license and changes it to a non IB-FLEX member. For an IB-FLEX appliance, it checks for an **Flex Grid Activation** Grid-wide license before node registration.

IB-FLEX members can join the Grid through the MGMT interface when Software ADP is enabled. You can configure an IB-FLEX appliance to function as a Grid Master or a member. To enable reporting for a Grid member that is running Software ADP, you must configure the MGMT interface.

A non IB-FLEX appliance designated as a member requires either a Grid and/or vNIOS/NIOS licenses installed to join the Grid. Similarly, for a reporting appliance to join the Grid, you must install a Grid and/or vNIOS/NIOS licenses. You cannot assign pool licenses to an IB-FLEX appliance. IB-FLEX supports HA for appliances that are running Software ADP.

Infoblox supports elastic scaling on IB-FLEX members that use the **Flex Grid Activation** Grid-wide license. It also supports pre-provisioning for Software ADP on the supported platforms. You must add the new IB-FLEX model to the list of supported pre-provisioning hardware types, so that you can select it during the member pre-provisioning. To pre-provision a non IB-FLEX Grid member, you must have valid pool licenses and pre-provisioned those members in the Grid.



### Important

To set up a supported virtual appliance as an IB-FLEX, you must first define the hardware type of the virtual appliance as IB-FLEX before you configure it. Depending on the platform or environment in which you are installing IB-FLEX, you can define the **hardware\_type** parameter to **IB-FLEX** during the cloud-init process, or you can manually set the hardware type using the **set hardware-type** CLI command. For more information, see [set hardware-type](#).

## Limitations of IB-FLEX

- It is not compatible with the traditional node-based licensing and it supports capacity based licensing only.
- An IB-FLEX instance will not start if you do not configure the required minimum level of resources.
- The resources assigned to IB-FLEX for cores and memory must be equal to or exceed the minimum designated values for the platform. For more information about IB-FLEX platforms, see About IB-FLEX Instances and Platform Settings below.
- IB-FLEX does not support DNS64 on appliances running NIOS version 8.2.0.
- IB-FLEX on AWS does not support ADP and DCA features.
- To effectively use the IB-FLEX Grid Master Candidate, it is mandatory to install FLEX Grid Activation license / FLEX Grid Activation for Managed Services license, on the Grid Master.

## Installing IB-FLEX

Depending on your network environment, you can install IB-FLEX just like how you install other Infoblox virtual appliances. Before you deploy an IB-FLEX, ensure that you set the hardware type of the appliance to IB-FLEX. You can do so either through the cloud-init process during deployment or manually through the **set hardware-type** CLI command.

For more information about installing IB-FLEX in the VMware environment, see [Deploying vNIOS Appliances on VMware](#).

For information about installing IB-FLEX in the OpenStack environment, see [Deploying vNIOS for KVM in OpenStack Using Elastic Scaling](#).

### About IB-FLEX Instances and Platform Settings

An IB-FLEX instance supports capacity-based licensing only, but it is compatible with NIOS Grid Master that uses node-based licensing. You can upgrade an IB-FLEX instance from a low-end platform to a high-end platform by increasing the resource allocation of the virtual machine. An IB-FLEX instance selects the default internal settings for a respective instance platform based on the resource settings detected during the startup.

An IB-FLEX instance supports VMware ESXi with or without SR-IOV enabled and OpenStack with KVM both with or without SR-IOV. The table below provides information about the IB-FLEX platform resource specification:

#### IB-FLEX Platform Resource Specification

Resource Type	Allowed Range of Values	Recommended Value	Description
Virtual NUMA Nodes	1	1	Single virtual CPU socket
Disk Size	250 GB	250 GB	Fixed size virtual disk

The table below provides information about the IB-FLEX platform and various platform settings:

#### Total Resource Usage for Different Use Cases

Intended Use	Total CPU	Total Virtual Memory GB (Without Software ADP)	Total Virtual Memory GB (With Software ADP)	Database Object Count	Grid Master Capable
Small Authoritative DNS	4	8	10	100,000	No
Medium Authoritative DNS	8	16	22	600,000	Yes
Large Authoritative DNS	16	32	40	5,000,000	Yes
Recursive DNS (without acceleration)	6	14	18	200,000	Yes
Large Recursive DNS (without acceleration)	14	28	36	2,000,000	Yes
Small Grid Master	10	18	NA	1,000,000	Yes
Medium Grid Master	12	22	NA	2,000,000	Yes
Large Grid Master	16	32	NA	16,000,000	Yes
Small Recursive DNS (with acceleration)	10	12	20	100,000	No
Medium Recursive DNS (with acceleration)	16	20	28	100,000	No

Intended Use	Total CPU	Total Virtual Memory GB (Without Software ADP)	Total Virtual Memory GB (With Software ADP)	Database Object Count	Grid Master Capable
Large Recursive DNS (with acceleration)	26	30	38	100,000	No
Large Grid Master (with acceleration)	20	38	NA	16,000,000	Yes

Note the following about IB-FLEX:

- You cannot mark an IB-FLEX appliance as a Grid Master or Grid Master Candidate with resources that are intended for small authoritative DNS, small recursive DNS (with acceleration), medium recursive DNS (with acceleration), and large recursive DNS (with acceleration). For more information, see the *Total Resource Usage for Different Use Cases* table above.
- Infoblox recommends that you increase the memory to the following for IB-FLEX members to use certain features:
  - 16 GB, instead of the standard 14 GB, to use DNS analytics.
  - 20 GB, instead of the standard 18 GB, to use Threat analytics when RPZ is assigned to the IB-FLEX member.

Consider the following recommendations for VMs on IB-FLEX appliances deployed in KVM-based OpenStack environment and with virtual DNS Cache Acceleration (vDCA) and/or software Advanced DNS Protection (vADP) enabled:

- vDCA and vADP are computational intensive features. Therefore, you must ensure that the vCPUs in each VM running vDCA and/or vADP does not exceed the total physical CPUs/cores of the NUMA node on which it is deployed. Refer to the following example configurations from: [Red Hat](#), [VMware](#), and [Oracle](#).
- Ensure that each vCPU of the VM is assigned to a separate physical core on the host and that it does not share the same physical core to avoid significant performance degradation that may occur.
- For deployments using SRIOV, the number of RX queues on a VF is dependent on the underlying driver. Most of the drivers limit the number of queues to 1, 2, or 4. As such, a large IB-FLEX deployment with vDCA will not be fully utilized if only the LAN1 interface is used as only 1, 2, or 4 out of 8 vDCA cores will be assigned to each interface/VF. Other cores will mostly remain idle unless the traffic also arrives on the LAN2 interface. Refer to the example configuration from [Red Hat](#).

## Configuring DNS Cache Acceleration on IB-FLEX

When you enable virtual DNS cache acceleration on the IB-FLEX, the appliance acts as a high-speed DNS caching-only name server. This feature provides DNS cache acceleration support for recursive UDP DNS queries on the IB-FLEX. The DNS cache acceleration feature is bundled with the **Tiered licensing**. When you install this license, you are entitled to use the DNS cache acceleration feature on IB-FLEX.

IB-FLEX supports RPZ, but the response for RPZ queries are not cached by the DNS cache accelerator. Instead, these queries are bypassed to the host and you can configure cache expiry period for RPZ queries. Note that the maximum cache lifetime for DNS cache acceleration on IB-FLEX is set to 300 seconds if the RPZ license is installed.

You can also use Elastic Scaling to pre-provision DNS cache acceleration on IB-FLEX. IB-FLEX supports Intel x86\_64 systems with IOMMU, Hugepages processors, virtio-net, and Intel 82599 10 G NIC and SRIOV with Intel 82599 ethernet controllers for DNS cache acceleration.

You can configure DNS cache acceleration on IB-FLEX using the Grid Manager or API. To view accelerated cache details, you can either log in to Grid Manager, or use CLI commands, or Infoblox API. If the tiered license usage is exceeded then a message is displayed in the Grid Manager. A warning message is displayed on the Grid Manager, if the QPS is going over the threshold on these platforms based on Tiered license installed.

Infoblox supports Auto Scaling that contains OpenStack packages to automatically scale the required number of resources based on your application. For more information, refer to [Auto Scaling for Virtual DNS Cache Acceleration](#). For detailed information about configuring DNS Cache Acceleration, see [Configuring DNS Cache Acceleration](#).



## Reports for IB-FLEX

Infoblox supports a selected set of reports on IB-FLEX. To view all available reports, from the **Reporting** tab, select the **Dashboards** tab. The table below lists all the supported reports for IB-FLEX. For information about how to create and manage user-defined reports, see [Infoblox Reporting and Analytics](#).

### Supported Reports for IB-FLEX

DNS Reports	Security (DNS) Reports	System Reports
DNS Query Rate by Query Type	DNS Top RPZ Hits	SPLA Grid Licensing Features Enabled
DNS Query Rate by Member	DNS Top RPZ Hits by Client	CPU Utilization Trend
DNS Daily Query Rate by Member	DNS RPZ Hits Trend By Mitigation Action	Memory Utilization Trend
DNS Daily Peak Hour Query Rate by Member		
DNS Replies Trend		
DNS Cache Hit Rate Trend		
DNS Top Requested Domain Names		
DNS Top NXDOMAIN / NOERROR (no data)		
DNS Top Clients		
DNS Top Timed-Out Recursive Queries		
DNS Response Latency Trend		
DNS Top SERVFAIL Errors Sent		
DNS Top SERVFAIL Errors Received		
DNS Object Count Trend for Flex Grid License		
DNS Effective Peak Usage Trend for SPLA Grid License		

## Configuring DNS Cache Acceleration

The software-based DNS acceleration feature supports IB-FLEX and non-IB-FLEX (IB-2215, IB-2225, IB-V2215, IB-V2225, IB-4015, IB-4025, IB-V4015, and IB-V4025) platforms. When you enable the virtual DNS cache acceleration feature on IB-Flex and non IB-Flex appliances, it acts as a high-speed DNS caching-only name server. This feature provides DNS cache acceleration support for recursive UDP DNS queries.

The DNS cache acceleration feature is bundled with tiered licensing for IB-FLEX appliances and for non-IB-FLEX appliances it is based on the type of tiered license that is installed. Only the Tier 1 (unlimited QPS up to capability) license can be installed on IB-2215 and IB-V2225 appliances. When you install the license, you are entitled to use the DNS cache acceleration feature. For non-IB-FLEX appliances, the warning message is based on the tiered license that is installed, and the QPS is rate-limited which is based on the type of license installed. If the tiered license and the QPS exceed the threshold, a warning message is displayed. For more information on the Tiered licensing feature, see the *Features on the Software DNS Cache Acceleration Platforms* table below for features on the Software DNS cache acceleration platforms.

All the appliances support RPZ, but the response for RPZ queries are not cached by the DNS cache accelerator. Instead, these queries are bypassed to the host. You can configure the cache expiry period for RPZ queries. Note that the maximum cache lifetime for DNS cache acceleration is set to 300 seconds if the RPZ license is installed. However, for IB-Flex appliances, you must configure RPZ zones for a member.

You can also use Elastic Scaling to pre-provision DNS cache acceleration. These appliances support Intel x86\_64 systems with IOMMU, Hugepages processors, virtio-net, and Intel 82599 10 G NIC and SRIOV with Intel 82599 ethernet controllers for DNS cache acceleration.

You can configure DNS cache acceleration using the Grid Manager or API. To view accelerated cache details, you can either log in to Grid Manager, or use CLI commands, or Infoblox API. Infoblox supports Auto Scaling that contains OpenStack packages to automatically scale the required number of resources based on your application. For more information, refer to [Auto Scaling for Virtual DNS Cache Acceleration](#).

Associated characteristics of the supported appliance include the following:

- Cache delete through the Grid Manager, CLI, or Infoblox API. For more information, see [Clearing DNS Cache](#).
- ACL for IPv4 and IPv6.
- Sending SNMP traps for DNS cache acceleration service.
- SNMP queries for supported appliances.
- Fixed RRSET order for accelerated responses, for A and AAAA record types, for IPv4, and IPv6.
- Both non-accelerated recursive and authoritative DNS with Software ADP.

The following table lists the features that are either supported or not supported on the Software DNS cache acceleration platforms:

### Features on the Software DNS Cache Acceleration Platforms

Features	IB-FLEX	IB-2215	IB-2225	IB-v2215	IB-v2225	IB-4015	IB-4025	IB-v4015	IB-v4025
Tiered licensing	Licensing is based on the Flex Grid Activation license on the Grid. Note that the queries per second are limited by the number of CPUs for IB-FLEX.	IB-40x5 appliances support four tiers of DNS QPS and the QPS levels are enforced by rate limiting							
RPZ	Yes, the maximum cache lifetime for DNS cache acceleration is set to 300 seconds if RPZ zones are configured for the member.	Yes, the maximum cache lifetime for DNS cache acceleration is set to 300 seconds if the RPZ license is installed.							
Caching (A, AAAA, MX, CNAME, PTR)	Yes	Yes							
Do not cache (EDNS, TCP, Any, TSIG)	Yes	Yes							

Features	IB-FLEX	IB-2215	IB-2225	IB-v2215	IB-v2225	IB-4015	IB-4025	IB-v4015	IB-v4025
Caching over additional interfaces (v4, v6)	Yes	Yes							
Dump Acceleration Cache (CLI, GUI, PAPI)	Yes	Yes							
Clear Acceleration Cache (CLI, GUI, PAPI)	Yes	Yes							
Cache pre-fetch and cache refresh	Yes	Yes							
ACLs (Allow-queries/ Responses, Match-Clients/Destination, Blackhole)	Yes	Yes							
AAAA Filtering (Bypassed but support configuring)	Yes	Yes							
Fixed RRSET ordering	Yes	Yes							
DNS64	Yes	Yes							
DNS monitoring feature (netmon)	Yes	Yes							
DNS Query logging (BIND only)	Yes	Yes							
DNS Views	Yes, it supports up to six DNS views.	Yes, it supports up to six DNS views.							
Forward/Stub zones	Yes	Yes							
Unbound as DNS resolver	Yes, unbound is supported through the Flex Grid Activation license.	Yes, unbound is supported if the Dual Engine DNS license is installed.							
DNS cache acceleration related restrictions for configuration	Yes, for NIOS version 8.2.0, restrictions are enforced based on whether the DNS cache acceleration feature is enabled or disabled.	No							
Reporting	Yes, for more information Reports for IB-FLEX, see <a href="#">About IB-FLEX</a> .	Yes							

Features	IB-FLEX	IB-2215	IB-2225	IB-v2215	IB-v2225	IB-4015	IB-4025	IB-v4015	IB-v4025
VLAN	Yes	Yes							
DSCP	No, Infoblox does not support DSCP for virtual appliances.	Infoblox does not support DSCP for physical or virtual appliances only if DCA is enabled.							
Sort list	Yes	Yes							
Anycast (OSPF and BGP)	Yes	Yes							
BFD (Bidirectional Forwarding Detection)	Yes	Supported on all appliances							
HA Support	Yes, only for non-SRIOV.	Yes							
NIC Bonding	Yes	Yes							
Multiple-Interfaces on the same subnet	No	No							
IP Rate-limit and Response logging	No	No							
EDNS Client Subnet support	No	No							
NXDomain-redirection	Yes	Ye							
DNSSEC (Bypassed but support configuring)	Yes	Yes							
Debug enhancements	Yes	Yes							
SNMP Support for DCA service-related traps	Yes	Yes							
SNMP stats support for DNS QPS and CHR	Yes	Yes							
NX Mitigation	No	No							
NetFilter (Tracking tables)	No	Not supported on any appliance							

Features	IB-FLEX	IB-2215	IB-2225	IB-v2215	IB-v2225	IB-4015	IB-4025	IB-v4015	IB-v4025
Traffic-capture (All modes)	Yes, there is partial support. Note that tcpdump captures both queries and responses.	Yes, there is partial support. Note that tcpdump captures both queries and responses.							
No flush-mode support for DNS cache acceleration cache	Yes	Yes							
Per-interface UDP DNS cache acceleration response counters	Yes	Yes							
CLI commands	You can use the commands <code>set dns-accel</code> and <code>show dns-accel</code> to view and set DNS cache acceleration information. For more information, see <a href="#">CLI Commands</a> .	You can use the commands <code>set dns-accel</code> and <code>show dns-accel</code> to view and set DNS cache acceleration information. For more information, see <a href="#">CLI Commands</a> .							
DNS Query rewrite (Bypassed but supports configuring)	No	No							
Threat Protection	Supported on IB-FLEX platforms. Allows enabling Software ADP and DNS cache acceleration simultaneously on IB-FLEX platforms.	Supported on IB-FLEX platforms. Allows enabling Software ADP and DNS cache acceleration simultaneously.							
Subscriber Secure Policy	Yes	Yes							



#### Note

By default, all malformed packets are dropped early when the accelerated threat protection service is enabled.

### Viewing Accelerated Cache Details

When you view cached contents of the DNS accelerator through the Grid Manager, there might be a slight impact on the DNS query performance of the selected member.

To view accelerated cache from the Grid Manager:

1. From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *Member* checkbox. Choose **View** from the Toolbar, and then click **View Cache**.
2. Click **Yes** in the *View Acceleration Cache* dialog box.
3. The system displays a *File Download was Successful* message and the cache data is displayed in table format in a new browser tab or browser window.

## Limitations for Virtual DNS Cache Acceleration

- You cannot enable the DNS cache acceleration feature during a scheduled NIOS upgrade, but if you have already enabled this feature, it will function normally during the upgrade process.
- The appliance prompts for a reboot when you enable the DNS cache acceleration feature for the first time. You must accept it to start the service.
- You must disable the DNS cache acceleration feature and reboot the appliance manually to switch from virtual DNS cache acceleration to authoritative servers.
- The appliance prompts for a reboot when you enable virtual DNS cache acceleration and Software ADP on IB-FLEX, IB-22x5 and IB-40x5 platforms .
- DSCP is not supported if vDCA is enabled on IB-FLEX 22x5 and IB-40x5.
- DHCP license cannot be installed if the DCA license is installed and vice versa.
- DCA and Microsoft Management licenses cannot be installed and configured simultaneously.
- On all vNIOS appliances that support vDCA (virtual DNS Cache Acceleration) or vADP (virtual Advanced DNS Protection), you must run vDCA or vADP on a single virtual NUMA node. If the configuration of the virtual NUMA node and physical NUMA node are not the same, it may result in performance degradation.
- For virtual DCA appliances using virtio, it is recommended to increase the number of virtio queues to 2 for IB-v14x5, 4 for IB-v22x5, 4 for IB-v40x5, 2 for IB-FLEX Small, 4 for IB-FLEX Medium and 4 for IB-FLEX Large systems.

## Limitations for DNS Cache Acceleration in Subscriber Services Parental Control

Enabling DNS Cache Acceleration for subscriber services in the **Parental Control** tab has the following limitations:

- The DNS Cache Acceleration subscriber site features, query count logging and blocked and allowed list support are applicable on Virtual DNS Cache Acceleration, and the features are available on IB-4030, however they are not supported.
- The DNS Cache Acceleration subscriber site features, query count logging and blocked and allowed list support and retains only unknown bits and does not support unknown policies (AVP).
- DNS Cache Acceleration uses BIND to process the guests behind Customer Premises Equipment (CPEs).
- The appliance prompts for a reboot when there is a configuration change.
- DNSTAP is required for query count logging.
- DNS Cache Acceleration does not cache blocked domains from BIND as it only uses category information for resolved domains.
- At Virtual DNS Cache Acceleration, the subscriber has access only to the primary MSP IP address.
- DNS Cache Acceleration subscriber site feature supports only 16 additional blocking policies.
- Before blocking another opt in subscriber at DNS Cache Acceleration, an opt in subscriber must resolve a domain.
- Proxy-All replies comes from DNS Cache Acceleration as long as the client connection status to MSP is "connected." If the client connection status is "disconnected," the first few queries goes to BIND, and future requests comes from DNS Cache Acceleration. Note that, TCP idle connections are closed every 20 seconds by MSP.
- The query name for the subscriber allowed and blocked list must contain a known TLD (top-level domain) and, if there are any prefixes, must conclude with a '.'.
- Only domain names are supported by the subscriber allowed and blocked lists, the wildcards and services are not supported.

For Information on Upgrading Parental Control at DNS Cache Acceleration, see [Upgrading Parental Control at DNS Cache Acceleration](#).

## IB-FLEX Platform Settings for DNS Cache Acceleration

When you enable the DNS cache acceleration feature on IB-FLEX, ensure that it has enough CPU and memory to start the service, and that it does not contain any authoritative zones. Note that you cannot start the service if the total CPU is less than 8 cores or if memory is less than 12G. To start the service, see the number of mandatory resources mentioned in the [Total Resource Usage for Different Use Cases](#) table.

If the DNS cache acceleration feature is enabled on a pre-provisioned member and fails to start due to insufficient resources on the member, the DCA status is displayed as failed. If you disable DCA on a member with insufficient resources, the member is not displayed in the **DCA** -> **Members** tab.



#### Note

- Under certain circumstances, the DNS cache acceleration feature may not function normally when you perform a product restart. This happens due to increased resource allocation on the virtual machine and the appliance does not log any entries in the syslog. Infoblox recommends that you restart or reboot the system and free up server resources if you encounter this issue.
- Before enabling DNS Cache Acceleration or ADP on virtual platforms, ensure that the `ssse3`, `sse4_1`, and `sse4_2` CPU flags are set on the host server. For more information, see <https://help.ubuntu.com/lts/serverguide/DPDK.html.en>
- If you see the `"/usr/bin/fast-path.sh: error starting /usr/bin/fp-rte. Check logs for details"` error message in the `infoblox.log` file, ensure that the `ssse3`, `sse4_1`, and `sse4_2` flags are set for the VM.

## About Elastic Scaling

Elastic Scaling provides the capability to automatically pre-provision and deploy vNIOS appliances on-demand for IPAM (IP Address Management), DNS, and/or DHCP. Compared to standard appliance deployment and licensing management, you now have the flexibility to purchase multiple service and feature licenses and install them as dynamic licenses for future vNIOS or cloud deployments based on your evolving business needs and deploy them as needed. When you remove a vNIOS or cloud appliance from the Grid, the dynamic licenses on the appliance are automatically released and returned to the license pool on the Grid Master for reuse at a later time. Elastic scaling includes a full set of APIs for pre-provisioning, deployment, and de-provisioning vNIOS appliances, making it simple to add or remove DNS or DHCP capacity on-demand to meet changing infrastructure requirements, which is critical for realizing the benefits of dynamic Cloud environments. For information about the Infoblox cloud solution, see [Deploying Cloud Network Automation](#).

You can purchase licenses of any type, such as vNIOS, DNS, DHCP, Enterprise (formerly Grid) and Cloud Platform, deploy them as dynamic licenses on the Grid Master, and allocate them to vNIOS virtual appliances manually or automatically through Elastic Scaling.

All dynamic licenses are tied to a specific Grid Master. You must first obtain the `LPC_UID` (License Pool Container Unique ID) of the Grid Master before obtaining the licenses from Infoblox Technical Support. For more information about Obtaining Dynamic Licenses, see [Obtaining Dynamic Licenses](#).

When you use Elastic Scaling to pre-provision and launch vNIOS appliances, licenses are automatically installed on the newly spun-up appliances during the process as long as you have the correct vNIOS licenses in the license pools for the vNIOS models you plan to deploy. For example, you can install a CP-V1405 license on a CP-V1405 or CP-V805, but you cannot do so on a CP-V2205. If there are no licenses available in the pool for the specified appliance models, the Grid Master notifies you with an error message (for Cloud Platform Appliances, the API calls fail). Note that you cannot add a vNIOS Grid member when no dynamic licenses are available. When you spin down vNIOS appliances, its assigned licenses are released and returned to the NIOS license pool.

The Grid Master keeps track of dynamic licenses that are allocated to vNIOS members and adjusts the total number of available dynamic licenses for each feature and service. You can view the total number of dynamic licenses installed for each feature and service, the number of active and available licenses, their usage, and other related information in the **Grid** tab -> **Licenses** tab -> **Pool** tab of Grid Manager. For information about how to view dynamic licenses, see [Viewing Dynamic Licenses](#).

Infoblox supports elastic scaling for Software ADP profiles on the supported platforms and provides the following pool licenses: **Threat Protection (Software add-on)** and **Threat Protection Update**. For more information about Software ADP profiles, see [Configuring Threat Protection Profiles](#). Threat protection members use management port for IPv4 and IPv6 communication with the Grid. Infoblox supports cloud API calls for such members to join the Grid using MGMT port and VPN on the MGMT port. To know more about using the MGMT port, see [Using the MGMT Port](#).

You can also manually allocate and deallocate dynamic licenses as your business requirements evolve. For more information about Manually Allocating Dynamic Licenses and Manually Deallocating Dynamic Licenses, see [Managing Licenses](#).

## Using Elastic Scaling to Pre-provision and Launch vNIOs Members

You can use Elastic Scaling to deploy on-demand vNIOs virtual members in a Grid and in your cloud environment and pre-provision them to manage networks and zones. For information about how to utilize Elastic Scaling to provision Cloud Platform Appliances and join them to the Grid using cloud API calls, see [Sample Cloud API Requests for Elastic Scaling](#).

Complete the following tasks to pre-provision and launch vNIOs appliances:

1. Obtain dynamic licenses and install them on the Grid Master, as described in [Obtaining Dynamic Licenses](#). You can also use a temporary license to spin up a Grid Master VM.
2. Create an offline Grid member or HA pair, as described in [Adding a Single Member or Adding an HA Member](#), see [Adding Grid Members](#).
3. Pre-provision the offline Grid member you just created, as described in [Configuring Pre-Provisioned Members](#), see [Pre-Provisioning NIOS and vNIOs Appliances](#). Note that there are a few guidelines that you might want to review before pre-provisioning your Grid member. For more information, see [Guidelines for Pre-provisioning Offline Grid Members](#).  
Note that you must include the **Grid** and **vNIOs** provisional licenses for pre-provisioned vNIOs members in order to join them to the Grid.
4. Generate a token for the Grid member, as described in [Generating Tokens for Grid Members](#) below. Note for HA pairs, the appliance generates two tokens—one for each node of the HA pair.
5. Use cloud API calls to add network views, networks, or zones and then delegate the objects to the offline Grid member. For information about sample cloud API requests, see [Sample Cloud API Requests for Elastic Scaling](#).
6. Use API requests to join the Grid member to the Grid. For sample API requests, see [Sample Cloud API Requests for Elastic Scaling](#). If for any reasons the automated process of Elastic Scaling fails or if you are unable to send API calls, you can use CLI commands to join the Grid member to the Grid as a workaround. For more information, see [Using CLI Commands to Join Grid Members](#) below.
7. Verify the Grid member has successfully joined the Grid, as described in [Viewing Status](#).
8. Verify the dynamic licenses have been allocated correctly by viewing the license usage, as described in [Viewing Dynamic Licenses](#).

## Generating Tokens for Grid Members

Before you can automatically allocate dynamic licenses to a pre-provisioned member, you must request a one-time token from the Grid Master. This token allows the member to register and authenticate itself to the Grid Master before a specified date and time (the default is 60 minutes from the time you generate the token). When the token is not used after the expiration date and time, it becomes invalid and you must generate another token for the member. You can configure the token usage timeout so the appliance can send syslog messages to alert you about the unused token. For information about how to set the token usage timeout value, see [Configuring Token Usage Timeout](#) below.

Using a one-time token eliminates the need for the Grid Master credentials to be exposed to other Grid members and the CMP (Cloud Management Platform) in the case of cloud implementation. Note that only superusers can generate and view the token for a pre-provisioned Grid member.



### Note

You can use API calls as part of the automated deployment process to generate a token for the vNIOs Grid member before joining it to the Grid. For information about sample API requests you can use to generate a token, see [Sample Cloud API Requests for Elastic Scaling](#). As a workaround, you can also generate a token through Grid Manager.

To generate a token through Grid Manager, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Click the Action icon



- next to the vNIOS member and select **Generate Token** from the list.
3. In the *Your Permission Token* dialog, the appliance displays the token and the **Expiration Date** of the token. You must generate a new token for the member if the token is not used before the expiration date.



#### Note

Copy this token and paste it at the CLI when you use the set token on command to set the token and generate the token file.

## Configuring Token Usage Timeout

You can configure the appliance to send syslog messages to alert you about an unused token that has been generated for a pre-provisioned member. Depending on the timeout interval you configure, the appliance sends a syslog message for each timeout interval until the token expires.

To configure the token usage timeout value:

1. From the **Grid** tab -> **Grid Manager** tab, click **Grid Properties** -> **Edit** from the toolbar.
2. In the *Grid Properties* editor, select the **General** tab -> **Basic** tab and complete the following:
  - **Token usage timeout:** Enter the time interval (in minutes) for which the appliance sends a syslog message to alert you about the unused permission token for a pre-provisioned member. For example, if you enter 5 here, the appliance sends a syslog message every five minutes. The default is 10.
3. Save the configuration.

## Using CLI Commands to Join Grid Members



#### Note

If for any reasons the automated process of Elastic Scaling does not function properly, you can use CLI commands to join Grid members to the Grid as a workaround.

When using Elastic Scaling, ensure that you have generated a token for the member, as described in *Generating Tokens for Grid Members* above, before joining the member to the Grid.

To join the vNIOS member to the Grid:

1. Access the Infoblox CLI using an SSHv2 connection through an SSHv2 client. You can also access the CLI by connecting a serial cable directly from the console port of a management system to the console port on the appliance, and then using a terminal emulation program such as Hilgraeve Hyperterminal® (provided with Windows® operating systems) and launch a session. The connection settings are:
  - Bits per second: 9600
  - Stop bits: 1
  - Data bits: 8
  - Flow control: Xon/Xoff
  - Parity: None
2. Log in using the default user name and password *admin* and *infoblox*. User names and passwords are case-sensitive.
3. To change the network settings from the default, enter the **set network** command. Then enter information as prompted to change the IP address, netmask, and gateway for the LAN1 port.

```
Infoblox > set network
```

```
NOTICE: All HA configuration is performed from the GUI. This interface is used only to configure a standalone node or to join a grid.
```

```
Enter IPv4 address [Default: n.n.n.41]: <Enter the LAN1 port IP address>
```

```
Enter netmask: [Default: 255.255.255.0]: <Enter the LAN1 port netmask>
```

```
Enter gateway address [Default: n.n.n.1]: <Enter the gateway IP address>
```

```
NOTICE: Additional IPv6 interface can be configured only via GUI.
```

```
Become grid member? (y or n): n
```

Note that you must enter **n** to use Elastic Scaling. If you enter **y**, the member becomes a Grid member and you will not be able to set token and join the pre-provisioned member to the Grid.

4. Use the `set token on` command to set the member token, the Grid Master IP address and certificate to the token file. Following is an example:

```
Infoblox > set token on
```

```
Enter GM-IP [Current: not defined]: <Enter the Grid Master IP address>
```

```
Enter Token [Current: not defined]: Copy token from the Your Permission
```

```
Token dialog in Grid Manager.
```

```
New Token Settings:
```

```
GM-IP: 1.1.1.1
```

```
Token: b25lLnZpcnR1YWxfbm9kZSQx
```

```
Is this correct? (y or n): y
```

```
Do you want to download the certificate form GM and validate (y or n): y
```

```
Is this correct and valid (y or n): y
```

```
Are you sure to apply and save settings to file?: y
```

```
The token and certificate are saved.
```

5. To verify the token:

```
Infoblox > show token
```

The CLI displays the current token setting and certification information. Verify this information.

Note if there is incorrect information, use `set token off` to remove the token file.

6. Use the `set token join` command to register the Grid member and get licenses from the license pool before joining the member to the Grid. Once the member joins the Grid, the token become invalid—you can use the token only once.

```
Infoblox > set token join
```

```
Are you sure to start Member registration Client? (y or no): y Starting
```

```
Member registration Client...
```

```
Connecting...
```



#### Note

For HA pairs, repeat the CLI commands on both nodes.

## Using OpenStack cloud-init template to configure Grid Master and join Grid members

You can use the following OpenStack cloud-init template to configure an IB-V815 as a Grid Master:

```
#infoblox-config remote_console_enabled: y default_admin_password: infoblox
```

```
temp_license: nios IB-V815 dns dhcp enterprise
```

```
lan1:
```

```
v4_addr: 10.2.0.132
v4_netmask: 255.255.255.0
v4_gw: 10.2.0.1
```

```
mgmt:
```

```
v4_addr: 10.1.0.69
v4_netmask: 255.255.255.0
v4_gw: 10.1.0.1
```

You can use the following OpenStack cloud-init template to join an IB-V815 member to the Grid:

```
#infoblox-config remote_console_enabled: y default_admin_password: infoblox
temp_license: nios IB-V815 dns dhcp enterprise sw_tp tp_sub lan1:
```

```
v4_addr: 10.2.0.140
v4_netmask: 255.255.255.0
v4_gw: 10.2.0.1
```

```
mgmt:
```

```
v4_addr: 10.1.0.77
v4_netmask: 255.255.255.0
v4_gw: 10.1.0.1
```

```
gridmaster:
```

```
token: xqyv+gEcPiUp9ETdHqmS2VcPIHEd81/U ip_addr: 10.39.8.109
join_intf: mgmt
certificate:-----BEGIN
```

```
CERTIFICATE-----MIIDdzCCA18CEBgaTP/
```

```
XX2lAxDokwClJub4wDQYJKoZIhvcNAQEFBQAwejELMAkGA1UEBh
```

```
MCVVMxEzARBgNVBAGTCkNhbgG1mb3JuaWExEjAQBgNVBAcTCVN1bm55dmFsZTERMA8GA1UEChMISW5mb2
Jsb3gx
```

```
FDASBgNVBAwTC0VuZ2luZWVyaW5nMRkwFwYDVQQDEXB3d3cuaW5mb2Jsb3guY29tMB4XDTE3MDMwNTE0
NTE1M1
```

```
oXDTE4MDMwNTE0NTE1M1owejELMAkGA1UEBhMCVVMxEzARBgNVBAGTCkNhbgG1mb3JuaWExEjAQBgNVBA
cTCVN1
```

```
bm55dmFsZTERMA8GA1UEChMISW5mb2Jsb3gxFDASBgNVBAsTC0VuZ2luZWVyaW5nMRkwFwYDVQQDExB
3d3cuaW
5mb2Jsb3guY29tMIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEAsRf7VSVyYgRZCsdEgqU5m5
31Pk0H q0LZ5CqWrcyGiKDYrbByPGATWS0KcQ9opUMj7VF3vttX0oY/
f2pI80AKr0r8ADWh70fqXFDWFAYsxGmP0dkFTd NajI0reIrlYE0tF3FTB0ZiXixfTUsI0hX96xNMU/
0tHptloQxXz9+UoIf7ovFi6D0QBwj tBHmcVYhIJh0CfRUm
MsIZgCupKVfwXNo3BMQfyNKsePj fVvoxCWTXF+KfAv3JS00ARbwuAZiYcMl2rdKb+8vBq4+IaMwr83Qa
JV8cph
Ahyt5s7PebgS+GJLWzcIdUXSecDl3HEpJxLMnV0ko8ZByN5T4mywz6GQIDAQABMA0GCSqGSIb3DQEBBQ
UAA4IB
  AQCWYwLb8Z5usHU0HL2WgyMkAZW8PYsjQNlv/aI/0kEkiJsvZc5H72frgbTA+whnz/
CqsRu8Rd06VEi+3UqR7n
  +0wRwSL6gWmlVBLNP3BZfsTKn0Bhd89hzUrSGtK07xF/
kY2qUEb6LnJ91B1046h7LUJutmzSPK2w10yY295kLe
NhQgG35oMWgztc7II6V7ViTnkqzEPWxILV0W1odIAodG46eyc0Cu5NPRWpN/FRn9gzSvL03YilJ4d/
bii31s0S
BZumFP+Q5e0i7bcElTmmhy5gsweITpfybUrFZAhXNs09832Ej11Q3lVKL42IDsixTKIFwbG+cNM7b7zf
C00j81
  ----END CERTIFICATE
```

You can use the following OpenStack cloud-init template to join an IB-V1415 member to the Grid:

```
#infoblox-config remote_console_enabled: y default_admin_password: infoblox
temp_license: nios IB-V1415 dns dhcp enterprise sw_tp tp_sub
#temp_license: nios IB-FLEX

lan1:

  v4_addr: 10.2.0.28
  v4_netmask: 255.255.255.0
  v4_gw: 10.2.0.1

ha:

  v4_addr: 10.2.0.30
  v4_netmask: 255.255.255.0
  v4_gw: 10.2.0.1

mgmt:
```

v4\_addr: 10.1.0.29

v4\_netmask: 255.255.255.0

v4\_gw: 10.1.0.1

gridmaster:

token: 0rPidqD1Iau91adaIL7zl07sZb0qxuk1 ip\_addr: 10.39.52.19

join\_intf: mgmt

certificate: -----BEGIN

CERTIFICATE-----MIIDdzCCA18CEChqLtGPEL/

kEVjEE488HtkwDQYJKoZIhvcNAQEFBQAweijELMAkGA1UEBh

MCVVMxEzARBgNVBAGTCkNhbgLmb3JuaWExEjAQBgNVBAcTCVN1bm55dmFsZTERMA8GA1UEChMISW5mb2Jsb3gx

FDASBgNVBAwTC0VuZ2luZWVyaW5nMRkwFwYDVQQDEXB3d3cuaW5mb2Jsb3guY29tMB4XDTE3MDIyMjA5MDEyOV

oXDTE4MDIyMjA5MDEyOVoweijELMAkGA1UEBhMCVVMxEzARBgNVBAGTCkNhbgLmb3JuaWExEjAQBgNVBAcTCVN1

bm55dmFsZTERMA8GA1UEChMISW5mb2Jsb3gxFDASBgNVBAwTC0VuZ2luZWVyaW5nMRkwFwYDVQQDEXB3d3cuaW

5mb2Jsb3guY29tMIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEA02LEbIeAjRZhbQSSPRIMoeR6GZC

SftQV+DHPQAmvzPeJqaH8obCcRi6pfrPToxTKRCde7W87Tdy/

uurZVXbJNwdtW7xhfeLFvmdFuUGR+PIId7oJd

nd9qmBLmUUPRniQDkk5pM8+g+oLwjXPv2yn+zaad+LaZpXUslP7TSfVvIeo6t2lwsUUxyozUnGLN9Pm91u/k/pz

Cog2e+3y/F2WPYQzmAC5KU5vY8Rl8iX8z/03eHhnVFITSrk15xgE5IQtlJG5C/RksFt/b5gcAFqh/7yUhCPvW2

pd8/xw/caXsY2nFUC1b3jgUg+EfxpXE7EMD/thxqkhMNNK9G0hPrbVdQIDAQABMA0GCSqGSIb3DQEBBQUAA4IB

AQBiTz2cbvFUHIoQiLefSaf5Yv1fM6AyZ/

sjPlVjYa0DB0dn4n1iiIL0tibPML3v3SVd2suAFPLmZdf1XTqkaT rN8SLE0RR7fS/

7Nz7eibPLXWGgeY6se8Br9cLWm+1AP7ugAPvjSZxBn87Spz6BfZKQ7L1NKHeqfu0UDuUvv2r0

tdLbRSHhb0INmm20LlMmLwLxTCg/o7W2YaJa9lgyzz20oaZHGD1dLEP+mh2TsRyX/fxXYpwiAvmZ/

VkccLgC xcj/fU44hXlfFa+Ibz5sjYp1gExYfGFwUBDuf/

7ftrBNh90qcXzXncrQAebGBHhRYtsDpRnpWH+qGAzTdJXTm8

---END CERTIFICATE---

You can use the following OpenStack cloud-init template to join an IB-V825 member to the Grid:

```
#infoblox-config remote_console_enabled: y default_admin_password: infoblox
temp_license: nios IB-V825 dns dhcp enterprise
```

```
lan1:
```

```
v6_addr: 2620:10a:6000:2708::17
v6_cidr: 64
v6_gw: 2620:10a:6000:2708::1
```

```
mgmt:
```

```
v6_addr: 2620:10a:6000:2701::a
v6_cidr: 64
v6_gw: 2620:10a:6000:2701::1
```

```
gridmaster:
```

```
token: IDUxCCzc/o08PHUURVVTG2KoeSUsq0x0 ip_addr: 2620:10a:6000:2701::8
join_intf: mgmt
certificate: -----BEGIN
```

```
CERTIFICATE-----
```

```
MIIDdzCCA18CEDdxmmpWBgZpzPXFj01fzowDQYJKoZIhvcNAQEFBQAwejELMAkGA1UEBh
MCVVMxEzARBgNVBAGTCkNhbG1mb3JuaWExEjAQBGNVBAcTCVN1bm55dmFsZTERMA8GA1UEChMISW5mb2
Jsb3gx
FDASBgNVBAwTC0VuZ2luZWVyaW5nMRkwFwYDVQQDEXB3d3cuaW5mb2Jsb3guY29tMB4XDTE3MDExMTEz
NDY0OV
oXDTE4MDExMTEzNDY0OVowEjELMAkGA1UEBhMCVVMxEzARBgNVBAGTCkNhbG1mb3JuaWExEjAQBGNVBA
cTCVN1
bm55dmFsZTERMA8GA1UEChMISW5mb2Jsb3gxFDASBgNVBAwTC0VuZ2luZWVyaW5nMRkwFwYDVQQDEXB3
d3cuaW
5mb2Jsb3guY29tMIIBIjANBgkqhkiG9w0BAQEFAAOCAQ8AMIIBCgKCAQEArDd9+rSVV7zah8S/
zRSFPtmiE00X 6SLPbXWftI+5PdVyiZl+IcGr10Z09hoNGddZvXTyzCJCKI/
J5WA9+PzjJJnjRRjGaB8QL3Pq8Dpe24VVpaL92 vSkHARuKS9IjNZk20xGrPC0EMvb8/8H6Q/
Ym1Wmm8IocHXxL9syaSG6lyhJstsaNDV/J1U7d0Qwmx/wJ00Zv2
pJshFKWGC8pnxH4IWSCPvKv1scJYmgUttVKzCz1AgQ6+qEbAMJXkfAF39Hak/
gKwLArENOheQVZg0lbZV6fhI1
etmXNsR84wzELT7h2Xe4i6Dd01g287MCZAaLXzqDSGAhfVKcBkaCvs4wIDAQABMA0GCSqGSIb3DQEBBQ
UAA4IB
```

```
AQACqr/Sjo8e07Vqp/gUhzwrV27NISHpI0VXp0/  
j2Vl4JZ3kkPAIDgcmT9flHr6QLJc2KsU2WVyt8XPB0XYWes jEJ4m468NVwGDkveDCnJ5le7/  
oYub3aKOYN8/Bkd5hju/GNmcKXybx8yPjw9hnXfG1sPT6H9UaMpxHx2cVH9se EvLNbxIF2hVg/  
yX0kE+YOQP892up9IANVKjSFCsQEKZ6os961IZjzY/MQYr4aWoP1KfU825chZ7BqCCDQMj0Vx  
CX2pHKzuFoCYB8a3/Tt0znlm/7u1RuHftqHAKLXeabLmxMJBW/  
5ZoX0RSjbr40vcekws2e7MuklnCMuSlJA2uL  
---END CERTIFICATE---
```

To configure an IB-FLEX Grid Master using the **Flex Grid Activation** license, you can use the following OpenStack cloud-init template:

```
#infoblox-config  
  
remote_console_enabled: y  
hardware_type: IB-FLEX  
temp_license: flex_grid  
lan1:  
    v4_addr: 10.39.51.33  
    v4_netmask: 255.255.255.0  
    v4_gw: 10.39.51.1  
  
mgmt:  
    v4_addr: 10.39.50.22  
    v4_netmask: 255.255.255.0  
    v4_gw: 10.39.50.1  
  
lan2:  
    nic_bonding_enabled: Y  
    bonding_failback_interface: lan1  
  
mac:  
    mgmt: fa:16:3e:14:3a:ae  
    lan1: fa:16:3e:01:29:0b  
    ha: fa:16:3e:25:43:8a  
    lan2: fa:16:3e:8e:26:4c
```

## Managing the Order of Match Lists

When you configure certain DNS and DHCP functions, you can create match lists that the appliance uses to filter specific IP addresses for specific operations. For example, you can create a DNS blackhole list for including and excluding DNS traffic to certain IP addresses, configure a list of IP addresses for allowing and denying DDNS updates, or define a Match Destinations list that identifies destination addresses and TSIG keys that are allowed access to a DNS view. The appliance matches rules in these lists from top to bottom. Rules at the top always take precedence over those at the

bottom. Therefore, ensure that you put the most specific rules at the top of the list, and then put the more general rules at the bottom. For example, when you add network 10.10.0.0/24 to a DNS blackhole list, all 256 IP addresses in the network are put on the blackhole list. To allow DNS traffic to the specific IP addresses 10.10.0.55 and 10.10.0.88, you must add these two addresses at the top of the blackhole list before the network address 10.10.0.0/24, and then set their permissions to "Exclude." The same applies when you set up the list of clients for DDNS updates. If you want to deny DDNS updates from a specific client (10.0.0.99) and allow DDNS updates from all other clients in the 10.0.0.0/24 network, you must put 10.0.0.99 at the top of the list and configure the appliance to deny DDNS updates from this client. You then add network 10.0.0.0/24 for allowing DDNS updates from all other clients at the bottom of the list.

## Managing NIOS Appliances

To reboot and shut down a NIOS appliance, you can use Grid Manager or the Infoblox CLI. To reset a NIOS appliance, you must use the Infoblox CLI. You can also restart a Grid member or restart a GUI session through Grid Manager. You can also force an HA failover using the GUI.

### Restarting a Grid Member

You can restart a single Grid member or an HA pair on the NIOS appliance. When you restart a Grid member, the appliance logs the task in the audit log.

To restart a single Grid member:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox.
2. Expand the Toolbar and click **Control** -> **Restart**.
3. In the *Restart Product on Member* dialog box, click **OK** to restart the Grid member.

### Rebooting a NIOS Appliance

You can reboot a single NIOS appliance, a single node in an HA pair, or both nodes in an HA pair. To reboot a single NIOS appliance or one or both nodes in an HA pair:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox.
2. Expand the Toolbar and click **Control** -> **Reboot**.  
For an HA pair, choose whether to boot one node (and which one) or both nodes, and then click **OK**. Depending on the browser you use, Grid Manager may display a dialog box that indicates the system is unavailable during a restart or reboot.

To reboot a single NIOS appliance using the CLI:

1. Log in to the Infoblox CLI using a superuser account for the NIOS appliance that you intend to reboot.
2. Enter the following CLI command: **reboot**

### Shutting Down a NIOS Appliance

Under normal circumstances, you do not need to turn off or shut down a NIOS appliance. It is designed to operate continuously. However, if you want to turn off a NIOS appliance, use the GUI or the CLI to shut down the appliance, instead of turning it off by using the power switch.



#### Note

If there is a disruption in power when the NIOS appliance is operating, the NIOS appliance automatically reboots itself when power is restored.

To shut down a NIOS appliance:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox.



2. Expand the Toolbar and click **Control** -> **Shutdown**.  
For an HA pair, choose whether to shut down one node (and which one) or both nodes, and then click **OK**.  
The NIOS appliance shuts down. The fans might continue to operate until the appliance cools down.

To shut down a NIOS appliance using the CLI:

1. Log in to the Infoblox CLI using a superuser account.
2. Enter the following CLI command: `shutdown`

### Forcing an HA Failover

If you want to change which node in an HA pair is active and which is passive, you can force a failover to occur. Within five seconds after initiating a failover, the previously passive node becomes active and assumes ownership of the VIP address. Note that a forced failover causes a temporary service disruption. To proceed with the forced failover, click **OK**. The appliance logs this task in the audit log.

To restart a single Grid member:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox.  
Note the Grid member that you select must be an HA pair.
2. Expand the Toolbar and click **Control** -> **Force HA Failover**.
3. In the *For HA Failover* dialog box, click **OK** to change the node in an HA pair.

### Restarting GUI Sessions

You can restart a GUI session through Grid Manager. When you restart GUI on a specific member, the appliance logs off all other GUI sessions that are currently running even though they are opened in different browsers or another system. The appliance logs this task in the audit log.

To restart GUI sessions on a Grid member:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox.
2. Expand the Toolbar and click **Control** -> **Restart GUI**.
3. In the *Restart GUI* dialog box, click **OK** to restart the NIOS appliance.

### Resetting a NIOS Appliance

There are three ways to reset a NIOS appliance:

- Resetting the Database
- Resetting a NIOS Appliance to Factory Settings
- Resetting the NIOS Appliance to Factory Settings and Removing Licenses You can perform these functions only through the CLI.

#### Resetting the Database

You can reset the database if you lose the administrator account and password or if you want to clear the database but preserve the log files to diagnose a problem. This function removes the configuration files, and the DNS and DHCP data from the appliance database. During this procedure, you are given the option to preserve the network settings of the appliance, which are the IP address and subnet mask, the IP address of the gateway, the host name, and the remote access setting.

To reset the database:

1. Log in to the Infoblox CLI using a superuser account.
2. Enter the following CLI command: `reset database`

The appliance then displays a message similar to the following:

```
The following network settings can be restored after reset: IP Address:
```

```
10.1.1.10
```

```
Subnet Mask: 255.255.255.0
```

```
Gateway: 10.1.1.1
```

```
Host Name: ns1.corpxyz.com Remote Console Access: true
```

```
The entire database will be erased.
```

```
Do you wish to preserve basic network settings? (y or n)
```

3. Press the **Y** key to preserve the network settings or the **N** key to return the network settings to their default values (192.168.1.2, 255.255.255.0, 192.168.1.1).

## Resetting a NIOS Appliance to Factory Settings

You can reset a NIOS appliance to its original factory settings. This removes the database, network settings, logs, and configuration files. Then, it reboots with its factory settings, which are the default user name and password, and default network settings. When you perform this procedure, the appliance does not give you the option to preserve your network settings.



### Note

If you have previously imported HTTPS certificates, the appliance regenerates the certificates and replaces them.

To reset the NIOS appliance to its factory settings:

1. Log in to the Infoblox CLI using a superuser account.
2. Enter the following CLI command: **reset all**

## Resetting the NIOS Appliance to Factory Settings and Removing Licenses

You can also reset a NIOS appliance to its original factory settings and remove all the licenses installed on the appliance. This removes the database, network settings, logs, configuration files, and licenses. The appliance then reboots with its factory settings, which are the default user name and password, and default network settings.



### Note

If you have previously imported HTTPS certificates, the NIOS appliance regenerates the certificates and replaces them.

To reset the NIOS appliance to its factory settings and remove all its licenses:

1. Log in to the Infoblox CLI using a superuser account.
2. Enter the following CLI command: **reset all licenses**.

## Managing the Disk Subsystem on the Infoblox-4010

Among its many features, the Infoblox-4010 use a RAID (Redundant Array of Independent Disks) 10 array to provide the optimum mix of high database performance and redundant data storage with recovery features in the event of disk failures. The disk array is completely self managed. There are no maintenance or special procedures required to service the disk subsystem.

---


Caution: Never remove more than one disk at a time from the array. Removing two or more disks at once can cause an array failure and result in an unrecoverable condition. You should replace only one disk at a time, using a replacement disk from Infoblox. For information, see [Replacing a Failed Disk Drive](#) below.

---

## About RAID 10

RAID 10 (or sometimes called RAID 1+0) uses a minimum of four disk drives to create a RAID 0 array from two RAID 1 arrays, as illustrated in RAID 10 Array Configuration below. It uses mirroring and striping to form a stripe of mirrored subsets. This means that the array combines—or stripes—multiple disk drives, creating a single logical volume (RAID 0). RAID 10 combines the high performance of RAID 0 and the high fault tolerance of RAID 1. Striping disk drives improves database write performance over a single disk drive for large databases. The disks are also mirrored (RAID 1), so that each disk in the logical volume is fully redundant.

### RAID 10 Array Configuration



When evaluating a fault on the Infoblox-4010, it is best to think of the disk subsystem as a single, integrated unit with four components, rather than four independent disk drives. For information, see [Evaluating the Status of the Disk Subsystem](#) below.

## Evaluating the Status of the Disk Subsystem

You can monitor the disk subsystem through the Infoblox Grid Manager GUI, the scrolling front panel LCD display, and four front panel LEDs next to the disk drives. In addition, you can monitor the disk status by using the CLI command `show hardware_status`. The following example displays the status of an Infoblox-4010 using the command:

```
Infoblox > show hardware_status
POWER:Power OK
Fan1:7258 RPM
Fan2:6887 RPM
Fan3:7258 RPM
CPU1_TEMP: +20.0 C
CPU2_TEMP: +24.0 C
SYS_TEMP: +35 C
```

RAID\_ARRAY: OPTIMAL




RAID\_BATTERY: OK READY Yes 103 HOURS

The *Detailed Status* panel provides a detailed status report on the appliance and service operations. To see a detailed status report:

- From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox -> Detailed Status icon in the table toolbar.

After displaying the *Detailed Status* panel, you can view the status of the selected Grid member. For more information on the *Detailed Status* panel, see [Status Dashboard](#).

The RAID icons indicate the status of the RAID array on the Infoblox-4010.

Icon	Color	Meaning
	Green	The RAID array is in an optimal state.
	Yellow	A new disk was inserted and the RAID array is rebuilding.
	Red	The RAID array is degraded. At least one disk is not functioning properly. The GUI lists the disks that are online. Replace only the disks that are offline.

The appliance also displays the type of each disk. In the event of a disk failure, you must replace the failed disk with one that is qualified and shipped from Infoblox and has the same disk type as the rest of the disks in the array.

Infoblox-4010 uses only the IB-Type 3 disk type. All disk drives in the array must have the same disk type for the array to function properly. You can have either IB-Type 1, IB-Type 2, or IB-Type-3, but you cannot mix both in the array. When you have a mismatched disk in the array, you must promptly replace the disk with a replacement disk from Infoblox to avoid operational issues.

### Disk Drive Front Panel LEDs

The disk drives of the Infoblox-4010 are located on the appliance front panel. To the right of each drive, two LEDs display connection and activity status. The below table lists the disk drive LED combinations and the states they represent.

#### *Infoblox-4010 Disk Drive LED Combinations*

Online/Activity LED (Green)	Fault/UID LED (Amber/Blue)	Description
On, off, or blinking	Alternating amber and blue	The drive has failed, or it has received a predictive failure alert; it also has been selected by a management application.
On, off, or blinking	Steadily blue	The drive is operating normally.
On	Amber, blinking regularly (1 Hz)	The drive has received a predictive failure alert. Replace the drive as soon as possible.
On	Off	The drive is online but it is not currently active.
Blinking regularly (1Hz)	Off	Do not remove the drive. The drive is rebuilding. Removing the drive may terminate the current operation and cause data loss.

Online/Activity LED (Green)	Fault/UID LED (Amber/Blue)	Description
Blinking irregularly	Amber, blinking regularly (1 Hz)	The drive is active, but it has received a predictive failure alert. Replace the drive as soon as possible.
Blinking irregularly	Off	The drive is active and operating normally.
Off	Steadily amber	A critical fault condition has been identified for this drive, and the controller has placed it offline. Replace the drive as soon as possible.
Off	Amber, blinking regularly (1 Hz)	The drive has received a predictive failure alert. Replace the drive as soon as possible.
Off	Off	The drive is offline, a spare, or not configured as part of an array.

## Replacing a Failed Disk Drive

The Infoblox-4010 is designed to provide continuous operation in the event of a failed disk. Replace an original RAID disk only when there is a disk failure. Hot-swapping a disk drive is a simple process that does not require issuing commands or a GUI operation.

When you replace a failed disk, you must replace it with an Infoblox supplied disk. To ensure that you receive the correct replacement disk, report the disk type or part number of the failed disk. The appliance displays the disk type in the *Detailed Status* panel, and the Infoblox part number is printed on the disk. Installing disks that are not qualified and shipped from Infoblox could cause failures in the appliance.

To replace a disk drive, follow this procedure:

1. Identify and verify the failed drive via the Grid Manager, front panel LCD, or CLI.
2. Make sure you have identified the correct drive.  
Note: Do not remove a correctly functioning drive.
3. Push in the latch for the drive and pull the release lever out towards you.
4. When the drive disengages, wait about 30 seconds for the disk to completely stop spinning.
5. Slide it out of the slot.

Replacement drives are shipped as a complete unit, ready to insert into the appliance. There is no preparation required. To install a replacement drive, follow this procedure:

1. Insert the replacement drive into the drive bay slot.
2. Gently slide the drive into place. When you feel the release lever engage, continue applying gentle pressure to the drive while pushing the release lever towards the appliance.
3. The release lever locks into place and the LED next to the disk drive lights up. Note that if the alarm buzzer is sounding, it automatically turns off about 20 seconds after the drive is inserted.
4. The disk drive automatically goes into rebuild mode.

## Disk Array Guidelines

Infoblox has designed the disk array to be completely self managed. There are no maintenance procedures required for a normally functioning disk array. Mishandling the disk array can cause an unrecoverable error and result in a failed appliance. Infoblox highly recommends that you observe the following guidelines:

- Remove only one disk at a time. Do not remove two or more disks from the appliance at the same time. Removing two or more disks at the same time might result in an appliance failure and require an RMA of the appliance. This rule applies to both powered and powered down appliances.
- If the status of the array is degraded, remove the failed or failing disk drive only. Do not remove an optimally functioning drive.
- If your acceptance procedure requires a test of the RAID hot swap feature, remove only one disk drive at a time. You can remove a second disk only after you replace the first disk and the array completes its rebuilding process.

- Do not remove a disk drive if the array is rebuilding. This could result in an appliance failure. Verify the status of the array before removing a disk drive.
- Use the following procedure to remove a spinning disk:
  - a. Unlatch and pull the disk about two cm (one inch) to disengage contact.
  - b. Wait about 30 seconds for the disk to completely stop spinning.
  - c. Remove the disk and handle it with care. Do not drop the disk or ship it loosely in a carton.
- You can hot swap a drive while the appliance remains in production.
- There are some conditions that may require powering down the appliance to replace a failed unit. This normally happens if the RAID controller detects an error that could damage the array. If you insert a replacement drive into a live array and the controller doesn't recognize the drive, power down the appliance.
- If you inadvertently remove the wrong disk drive, do not immediately remove the disk drive that you originally intended to remove. Verify the status of the array and replace the disk drive that you removed earlier before removing another drive. Removing a second drive could render the appliance inoperable.
- Older appliances have an audio alarm buzzer that sounds if a drive fails. The alarm automatically stops about 20 seconds after a functional disk has been inserted into the array.
- All disks in the RAID array should have the same disk type for the array to function properly.
- In the unlikely event that two disk drives fail simultaneously and the appliance is still operational, remove and replace the failed disk drives one at a time.
- Rebuild time depends on a number of factors, such as the system load and Grid replication activities. On very busy appliances (over 90% utilization), the disk rebuild process can take as long as 40 hours. On a Grid Master serving a very large Grid, expect the rebuild process to take at least 24 hours.
- Replace a failed or mismatched disk only with a replacement disk shipped from Infoblox. When you request a replacement disk, report the disk type displayed in the *Detailed Status* panel of the GUI or the Infoblox part number on the disk.

## Restarting Services

Whenever you make a change such as adding a zone or a network, Grid Manager notifies you that a service restart is required. You can enable the appliance to display the **Restart Banner** whenever it requires a service restart and prompt the administrator to input the user name before restarting the services. For information about how to enable the restart banner, see [Changing Grid Properties](#). To view all pending activities that are waiting for service restart, you can click **View Changes** in the restart banner at the top of Grid Manager.

The pending activities include additions, modifications, and deletions performed by all administrators. You can also view pending changes through the *Restart <Grid/ Member> Services* dialog box.

There are several ways to restart services on the affected Grid members. You can restart services at the Grid or member level, or you can restart services by groups, as described below:

- Restarting Grid Services
- Restarting Member Services
- Restarting Services by Groups

You can configure services restart settings such as restart timeout or delay as described in [Configuring Restart Settings](#) below.



### Note

When you make configuration changes for DNS or DHCP and the service is enabled on at least one Grid member, Grid Manager suggests a restart even if the service is disabled on the member affected by the change.

## User Permissions for Restarting Services

The following rules apply to superusers and limited-access users:

- You can cancel a schedule task that you create to restart a service. A superuser can cancel any scheduled restart task.

- Only superusers and administrators with read/write permission to all Grid members can schedule a Grid restart task.
- When a superuser schedules a Grid restart task, a limited-access user cannot schedule a member-level restart task.
- Limited-access users cannot cancel a superuser's scheduled tasks.
- Limited-access users cannot create or modify a scheduled restart task for a Grid member if a scheduled restart task for the member (created by another user) already exists.

The system writes every scheduled change action to the audit log as follows:

```
USER logon_id action service restart schedule 'schedule' on Grid (or member)
Grid name or member node id
```

For example:

```
USER jdoe insert service restart schedule '02/20/2007 01:30:00' on Grid
Infoblox USER jdoe deleted service restart schedule '02/22/2007 01:30:00' on
node id 3
```

For more information on the audit log, see [Using the Audit Log](#).

## Restarting Grid Services

Only superusers and administrators with read/write permission to all Grid members can schedule a Grid restart task. You can restart services at the Grid level either simultaneously or sequentially, and can also specify the restart time. If you schedule the restart at a specific date and time, the system schedules the restart at the specified time on each Grid member.

To restart services at the Grid level:

1. From the **Data Management** tab, select the **DHCP**, **DNS**, or **Grid** tab, or select the **Administration** tab, and then click **Restart Services** from the Toolbar.

The *Restart Grid Services* dialog appears.

1. You can specify whether the member restarts services when necessary or you can force it to restart services. Select one of the following in the [RestartGridServices](#) section:
  - **Ifneeded**: Select this to restart all active DNS and DHCP services if there are any changes requiring a service restart.
  - **Forceservicerestart**: Select this to force a service restart even if it is not needed. A forced restart may be delayed if there are pending restarts for the same service.
2. Select one of the following in the *RestartMethod* section:
  - **RestartAllRestartGroups**: Restarts the services in the affected Grid members in the order defined for restart groups.
  - **Simultaneouslyforallmembers**: Restarts the services on all of the members in a Grid at the same time.
  - **Sequentiallyforallmembers**: Restarts the services on each Grid member one after one.
  - **AffectedMembersandServices**: Click the Poll Members icon to display the affected members and services when the system restarts. Grid Manager displays the member names and one of the following for each service:
    - **YES**: The service is active and the system will restart the service upon execution of this task.
    - **NO**: The service will not restart unless the **Forceservicerestart** option is selected.
    - **DISABLED**: The service is currently disabled.
  - **ViewPendingChanges**: You can view the list of pending changes that will take effect when you restart the services. You can use filters to look for specific objects and view the following information for each pending activity:
    - **Timestamp**: The timestamp of the operation.
    - **Admin**: The admin user who performed the operation.
    - **Action**: The type of operation that was performed by an administrator. This can be **Created**, **Modified**, **Deleted**, **Called**, or **Message**.

- **ObjectType:** Displays the object type. For example, DNS View, Named ACL, IPv4 MAC Filter, Blacklist Rule, and so on.
  - **ObjectName:** The name of the object.
  - **Message:** Description of the activity.
3. To schedule a service restart task, click the Schedule icon at the top of the wizard. In the *ScheduleChange* panel, complete the following:
    - **Now:** Restarts services upon clicking **Restart**.
    - **Later:** Enter the following information to schedule all Grid members to restart services at a certain date and time:
      - **Date:** Enter a date in YYYY-MM-DD (year-month-day) format. The appliance displays today's date. You can also click the calendar icon to select a date from the calendar widget.
      - **Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. When you enter the time in a 24-hour format such as 23:00, Grid Manager displays 11:00:00 PM. You can also select a time from the drop-down list by clicking the time icon.
      - **TimeZone:** Select a time zone for the scheduled date and time from the drop-down list. This field displays the time zone of the browser that the admin uses to log in to Grid Manager.
  4. Click **Restart** to restart services immediately or click **ScheduleRestart** to create a scheduled restart task.



#### Note

When you restart services at the Grid level, the services are restarted only on those members for which you have permissions.

## Restarting Member Services

The member restart time always supersedes the Grid restart time. If the member restart time is later than the Grid restart time, then the member restarts services at its scheduled time. If the member restart time is before the Grid restart time, then the member restarts services at its scheduled restart time, and again during the Grid restart time.

To restart member services:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox.  
or  
From the **Grid** tab, select the **Grid Manager** tab, and then select a member checkbox.
2. Click **Restart Services** from the Toolbar.
3. You can specify whether the member should restart services when necessary or you can force it to restart services. Select one of the following in the *Restart Member Services* section:
  - **If needed:** Select this to restart all active DNS and DHCP services, if there are any changes requiring a service restart.
  - **Force service restart:** Select this to force a service restart even if it is not needed. A forced restart may be delayed if there are pending restarts for the same service.  
**Affected Services:** This table displays the affected services when the system restarts. It can display one of the following for each service:
    - **YES:** The service is active and the system will restart the service upon execution of this task.
    - **NO:** The service will not restart unless the **Force service restart** option is selected.
    - **DISABLED:** The service is currently disabled.**View Pending Changes:** You can view the list of pending changes that will take effect when you restart the services. You can use filters to look for specific objects and view the following information for each pending activity:
    - **Timestamp:** The timestamp of the operation.
    - **Admin:** The admin user who performed the operation.
    - **Action:** The type of operation that was performed by an administrator. This can be **Created, Modified, Deleted, Called,** or **Message**.
    - **Object Type:** Displays the object type. For example, DNS View, Named ACL, IPv4 MAC Filter, Blacklist Rule, and so on.
    - **Object Name:** The name of the object.
    - **Message:** Description of the activity.



4. To schedule a service restart task, click the Schedule icon at the top of the editor. In the *Schedule Change* panel, complete the following:
  - **Now:** Restarts services immediately.
  - **Later:** Enter the following information to schedule the member to restart services at a certain date and time:
    - **Date:** Enter a date in YYYY-MM-DD (year-month-day) format. The appliance displays today's date. You can also click the calendar icon to select a date from the calendar widget.
    - **Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
    - **Time Zone:** Select a time zone for the scheduled date and time from the drop-down list. This field displays the time zone of the browser that the admin uses to log in to Grid Manager.
5. Click **Restart** to restart services immediately or click **Schedule Restart** to create a scheduled restart task.

## Restarting Services by Groups

You can use the group restart feature for the DHCP and DNS services. You do this by adding Grid members to groups and defining the restart order. Each service has a separate set of groups. You can specify whether a group is restarted sequentially or simultaneously and can also create a recurring schedule for restarts.

To access restart groups:

1. Click **Grid** → **Grid Manager** → **DHCP** or **DNS**.
2. On the **Services** tab, click **Toggle Restart Groups View**.

You can arrange group restart order for each service independently. Members in a group can restart simultaneously or sequentially, depending on your choice. When members are started sequentially, they are started in alphabetical order. When you have several restart groups, they all restart sequentially in the order that you define. For information, see [Ordering Restart Groups](#) below.

For each service, there exists the Default restart group in Grid Manager: Default DHCP group and Default DNS group. All Grid members belong to the Default group until you add them to another custom restart group. By default, the members of the Default group are restarted sequentially. The Default group is always restarted in the last place in the restart sequence for a service.



### Note

You can manage restart groups only if you are a superuser. If you are a limited-access user, you can see the restart groups and their restart status only if these groups include members for which you have permissions. When you restart services by groups, the services are restarted only on those members for which you have permissions.

For more information on how to create and manage restart groups, see the following sections. For how to restart services by groups, see [Restarting Groups](#) below.

## Creating a Restart Group

You can add a member to a group if the member has a license for the corresponding service. Only members with the appropriate license are available for grouping. For pre-provisioned members, you can add them to groups, but the restart requests are not created for such members.

When you assign members to restart groups, Grid Manager checks if the restart may affect the service operation. If the restart with the current grouping configuration can lead to a service interruption, a warning message is displayed. For example, you cannot add two members from a DHCP failover association to the same group configured to restart its members simultaneously.

To create a restart group:

1. Click **Grid** → **Grid Manager** → **DHCP** or **DNS**.
2. On the **Services** tab, click **Toggle Restart Groups View**.
3. Click the **Add** icon.
4. Specify the general information for the group:
  - Name

- Comment
  - Restart order for group members: simultaneously or sequentially
- Note that for how to delay service restarts, see [Configuring Restart Settings](#) below.
5. Click **Next**.
  6. Add members by clicking the **Add** icon for each new member.
  7. Click **Next**.
  8. If you want to create a restart schedule for this group, select **Enable Restart Schedule** and specify the required parameters.  
Note that the restart schedule can run either once or on a recurring basis. It does not create scheduled tasks. If a schedule is configured for a restart group, you can still perform one-time restarts independently for Grid members or restart groups and create scheduled tasks for these restarts.
  9. Select the restart type for the services on the affected members:
    - **If needed**: Select this to restart services only if there are changes requiring a service restart.
    - **Force service restart**: Select this to force a service restart even if it is not needed. A forced restart may be delayed if there are pending restarts for the same service.
  10. Click **Next**.
  11. If necessary, specify the extensible attributes.
  12. Click **Save & Close**.

### Editing a Restart Group

You can edit a restart group at any time, even when the restart process is running. The changes do not affect the current restart process.

To edit a restart group:

1. In the restart groups list, select the checkbox of the group to edit.
2. Click the Edit icon.
3. Edit the group properties as described in [Creating a Restart Group](#) above.
4. Click **Save** or **Save & Close**.

### Deleting a Restart Group

You can delete any restart group, except the Default group. When you delete a group, all its members are automatically moved to the Default restart group. You can delete a restart group at any time, even when the restart process is running. The deletion does not affect the current restart process.

To delete a restart group:

1. In the restart groups list, select the checkbox of the group to delete.
2. Click the Delete icon.

### Ordering Restart Groups

You can modify the order in which multiple groups are restarted. You cannot change the order for the Default group, as it always restarts in the last position by default.

To order restart groups:

1. Go to the restart groups view for the DNS or DHCP service.
2. In the toolbar, click **Order Restart Groups**.
3. Reorder the groups by using the drag-and-drop operation or by clicking the up and down arrows.
4. Click **OK**.

### Restarting Groups

When you already have restart groups defined, you can restart services by specific groups at any time after you make configuration changes.

To restart services by groups:

1. Go to the restart groups view for the DNS or DHCP service.
2. Select the groups to restart.

3. In the toolbar, click **Restart → Restart Groups**.
4. In the Restart Group Confirmation window, select the restart type for the services on the affected members:
  - **If needed**: Select this to restart services only if there are changes requiring a service restart.
  - **Force service restart**: Select this to force a service restart even if it is not needed. A forced restart may be delayed if there are pending restarts for the same service.
5. To schedule a service restart, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, complete the following:
  - **Now**: Restarts services immediately.
  - **Later**: Enter the following information to schedule the member to restart services at a certain date and time:
    - **Date**: Enter a date in YYYY-MM-DD (year-month-day) format. The appliance displays today's date. You can also click the calendar icon to select a date from the calendar widget.
    - **Time**: Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
    - **Time Zone**: Select a time zone for the scheduled date and time from the drop-down list. This field displays the time zone of the browser that the admin uses to log in to Grid Manager.
6. Click **Restart** to restart services immediately or click **Schedule Restart** to schedule the restart.



#### Note

Normally, you cannot restart or schedule a restart of services when a scheduled full Grid upgrade is in progress. However, you can do so for the Default DNS or DHCP restart group. In this case, only the Grid Master is restarted. The other members of the group remain in the Timed Out status and are restarted after Grid Manager gets a response from them. For more information about the full Grid upgrades, see [Guidelines for Scheduling Full Upgrades](#).

## Configuring Restart Settings

To configure the service restart settings for the Grid:

1. Click **Grid→Grid Manager→DHCP** or **DNS** tab.
2. In the toolbar, click **Edit→Grid DHCP Properties** or **Grid DNS Properties**.
3. Click **Toggle Advanced Mode→Restart**.
4. If necessary, modify the service restart settings:
  - **Delay between Restart Groups**. This is the delay between the restarts of separate restart groups. The default is 10 seconds.



#### Note

If the delay time from the previous group restart has not elapsed yet, the new restart requests are displayed with the Pending restart status. To speed up the new restarts, you change the delay between restart groups to a smaller value at any time.

- **Member Restart Timeout**. This is the amount of time that the appliance waits for a response from a member. The default is 1 minute.
5. Optionally, select the checkbox **Apply restart requests to offline members when they connect**.
  6. Click **Save & Close**.

## Canceling a Scheduled Restart Task

Limited-access users can only cancel a scheduled task that they created. Superusers can cancel a scheduled task that any user created. You can cancel scheduled restart tasks from **Task Manager**. For information, see [Viewing Tasks](#). When you delete a scheduled restart task, the system cancels the scheduled task to restart services on the member or Grid and does not restart services. To cancel a scheduled restart, see [Canceling Scheduled Tasks](#). You can also turn off the automatic service restart feature for other scheduled tasks. For example, if you set a recurring

automatic restart for a restart group, you may not need the automatic service restart which should normally happen after a scheduled change of an object. You can do so from the Infoblox CLI using the `set scheduled` command. For information, see the section *set scheduled* in the *Infoblox CLI Guide*.

Alternatively, administrator users with the approver role can cancel the automatic service restart for individual tasks when approving them. If you are an approver, you can do so by deselecting the **Enable automatic service restart** option for a task in the **Task Manager**. For information, see [Viewing Tasks](#).

## Configuring the Orphan Mode

After you enable NTP on a Grid, the Grid members including Grid Manager can function as NTP servers to clients. For information on how to configure a NIOS appliance as an NTP server, see [NIOS Appliances as NTP Servers](#).



### Warning

Setting a proper stratum value is important for time and service synchronization among Grid members. Unless a special configuration is required, use the default values. In case you configure the values, keep the configuration as simple as possible.

The NIOS NTP service is configured with external NTP servers, and the NTP service stratum is derived from these external NTP servers. Sometimes, in the absence of the external NTP servers, NIOS Grid members, through the disconnected NTP server deliver the NTP service to the clients. Depending on your admin permissions, you need to configure the NTP service stratum values for Grid Manager and Grid members to facilitate the disconnected NTP service. The disconnected NTP service uses the service stratum values to help Grid Manager or members act as NTP servers. Once the external NTP server is accessible, the Grid connects to the external NTP server and receives the stratum values from the external NTP server. When the NTP service is in a disconnected state, it is referred to as the orphan mode.

You can configure stratum values for Grid Manager and Grid member at the Grid level, and for individual member at Grid member level using Grid Manager or CLI commands. For information on how to configure orphan mode stratum values, see [Defining NTP Orphan Mode](#) and for information about admin permissions, see [About Administrative Permissions](#).



### Note

The Grid member value configuration takes precedence over Grid Manager value configuration.

## Limitation of Orphan Mode

Note the following limitation of orphan mode:

- You can use the disconnected NTP service only if you have configured stratum values for NTP services on Grid Manager or a Grid member.
- When the member is offline and the Grid member and Grid Manager are at the default stratum, NIOS changes the stratum value of the local clock from 14 to 12.

## Guidelines for Setting Stratum Values

The following is a list of guidelines for setting stratum values:

- Either use the default NTP stratum values for Grid Manager as well as Grid members, or configure the NTP stratum values for Grid Manager as well as Grid members.
- To avoid the scenario where Grid Manager does not have internet access for a long period of time, add an external NTP server for the Grid member.

- With the default NTP stratum values, when the external NTP server is not reachable, Grid Manager sets the stratum value of the fallback local source to 12 even though the actual NTP stratum value is 1 greater than the fallback local source's stratum value, 13.
- When a Grid member synchronizes with Grid Manager, the NTP stratum value of the Grid member is 1 greater than the NTP stratum value of Grid Manager, 14. However, if Grid Manager is not suitable as an NTP clock source, then the Grid member uses its fallback local source's stratum value, 14, and the actual NTP stratum value of the Grid member is 1 greater than the fallback local source's stratum value, 15.
- When you configure the NTP stratum values, Grid Manager uses the fallback local source's stratum value of 1 less than the set NTP stratum value. For example, if the NTP stratum value is set to 8, the fallback local source's stratum value of Grid Manager is set to 7; if the external NTP servers are not reachable, then the NTP stratum value of Grid Manager becomes 8.
- When the external NTP servers are either not reachable or not configured, Grid Manager uses its local source and sets the NTP stratum value. In this scenario, however, the following applies to Grid members:
  - If a Grid member is configured to synchronize with only Grid Manager, then
    - If Grid Manager is reachable, then the Grid member's stratum value is 1 greater than the NTP stratum value of Grid Manager.
    - If Grid Manager is not reachable, then the Grid member's stratum value is set to the member's stratum value configured either at the Grid level or member level.
  - If the Grid member is configured to synchronize with only a set of external NTP servers, then the Grid member's stratum value is set to 1 greater than the NTP stratum value of the server that has the best quality and the lowest stratum value. However, if none of the external NTP servers can be reached, then the Grid member's stratum value is set to the member's stratum value configured either at the Grid level or member level.
  - If the Grid member is configured to synchronize with external NTP servers as well as Grid Manager, then
    - If the NTP servers as well as Grid Manager are reachable, then the Grid member synchronizes with the best available NTP source according to its reachability, stratum values, jitter, and offset, by setting its stratum value to 1 greater than the NTP stratum value of the best source.
    - If the external NTP servers are not reachable but Grid Manager is reachable, then the NTP stratum value is set to 1 greater than the NTP stratum value of Grid Manager.
    - If the NTP servers and Grid Manager are not reachable, then the Grid member synchronizes with the local clock and the stratum value is set to the member's stratum value configured either at the Grid level or member level.

## File Distribution Services

This section describes the file distribution services on the NIOS appliance. It contains the following topics:

- [About File Distribution Services](#)
- [File Distribution Storage](#)
- [Managing File Distribution Services](#)
- [Managing Directories](#)
- [Managing Files](#)
- [Viewing Files](#)
- [Managing Users](#)

### About File Distribution Services

You can upload files to the NIOS appliance and to individual Grid members using TFTP, HTTP, and FTP, clients. You can also upload files using the Grid Manager web interface or the API. Using access control lists, you can specify which network devices can upload files or retrieve files. You can use the **Group Results** function for file distribution services (TFTP, FTP, and HTTP) to group members by extensible attributes that contain the same values. For information about how to group members by extensible attributes, see [Adding Grid Members](#).

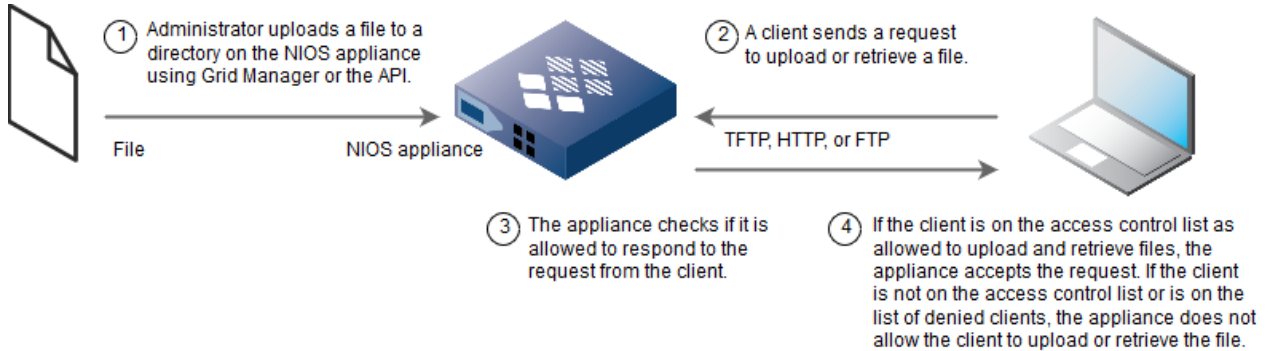


#### Note

File distribution services using TFTP, HTTP, and FTP is not supported by IPv6-only appliances.

Network devices, such as VoIP phones, can use the DHCP services on the appliance for IP address assignments and use the file distribution services for IP device configuration downloads. Downloads can be accomplished with TFTP, HTTP, or FTP.

#### Uploading and retrieving files



File uploads and downloads by FTP and TFTP file distribution clients are logged in the syslog under the **Administration->Logs** tabs. The logs store the following information:

- Client IP
- Date and Time
- Event type
- File(s) downloaded and/or uploaded

## File Distribution Storage

This section describes the storage capacity for file distribution, and how to manage file distribution storage settings. Maximum storage space allowed for all file distribution services on a Grid is equal to the storage space allowed on the Grid member with the smallest amount of space allowed.

Maximum storage space is shown in the below table.

#### Maximum Storage Space by Platform Type

Member Type	Description	Max Limit
INFOBLOX_MEMBER	All Infoblox appliances	10GB
VNIOS_MEMBER	All virtual appliances (VMWare)	5GB
VM_MEMBER	Virtual IPAM member (IPAM Free Ware)	1GB

## Usage Threshold Alerts

An SNMP trap generates an alarm message when a member nears storage capacity. The default threshold value is 90%, and the default reset value is 70%. If email notification is enabled, NIOS sends an email when either of these thresholds are reached.

- When the Grid member storage capacity reaches 100%, the SNMP trap generates a "High Usage" message. For information on how to modify the threshold values, see [Configuring SNMP](#).

- File distribution clients will fail to PUT files if the file is large enough that it will put the member over the storage limit.

## Modifying File Distribution Storage Limits

1. From the **Data Management** tab, select the **File Distribution** tab, and then click **Grid File Distribution Properties** from the Toolbar.
2. In the *Grid File Distribution Properties* editor, complete the following:
  - **Storage Limit (MB)**: Enter the maximum storage space in megabytes.
  - **Include files and directories in system backup**: This is selected by default to ensure that the appliance includes the uploaded files in the backup. You can clear this checkbox to improve the backup performance if you have stored these files separately.
3. Save the configuration and click **Restart** if it appears at the top of the screen.  
Note that to avoid data loss, after you change the storage limit FD services will be disrupted briefly and will take some time to resume. Wait until the File Distribution services are running again on the members before you upload any files.

## Managing File Distribution Services

This section describes how to configure file distribution services such as TFTP, FTP and HTTP. This section also describes how to configure access control lists which determine which clients are granted access to the service, and which clients are denied access to the service.

### Configuring the TFTP Service

The TFTP file distribution service is disabled on the appliance by default. To allow file distribution access using TFTP, you must specify the clients that are allowed to use the service and then enable the service on the appliance. If you do not specify this information or enable the service, the appliance denies access to all clients. The appliance provides read-only access to the files.

The TFTP service is supported only on LAN1 and MGMT interfaces. For more information, see [Configuring Ethernet Ports](#).

To configure the TFTP file distribution service on a member:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.
2. In the *Member File Distribution Properties* editor, select the **TFTP** tab, and then complete the following:
  - **Listen on Port**: Enter the number of the port on which the appliance receives TFTP file distribution requests.  
The default is port 69.
  - **Allow file transfers from**: Configure the appliance to grant or deny permissions to TFTP file distribution requests from clients, as described in [Configuring Access Control for File Distribution](#) below.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

After you configure the TFTP service, you must enable the service to allow file distribution access. For information, see [Starting and Stopping File Distribution Services](#) below.

### Configuring the FTP Service

The FTP file distribution service is disabled on the appliance by default. To allow file distribution access using FTP, you must create at least one user (see [Managing Users](#)), specify the clients that are allowed to use the service, and then enable the FTP service on the appliance. If you do not specify this information or enable the service, the appliance denies access to all clients. User creation is not necessary to access the FTP service if anonymous is enabled at Grid level. The appliance provides read-only access to the files.

To configure the FTP file distribution service on a member:



1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.
2. In the *Member File Distribution Properties* editor, select the **FTP** tab, and then complete the following:
  - **Listen on Port:** Enter the number of the port on which the appliance receives FTP file distribution requests. The default is port 21.
  - **Login Banner:** Enter your own login banner text that appears after you establish an FTP connection or use the default (**Restricted Access Only**).
  - **FTP Passive Mode:** By default, this is selected to enable FTP in passive mode; otherwise, it is in active mode. An FTP connection between a client and server can be in active or passive mode. In active mode, the server initiates the data connection. In passive mode, the client initiates the data connection. Depending on your firewall policy, firewalls can block active mode connections. There is no firewall filtering in passive mode.
  - **FTP File Listing:** Select this to allow users to list files and subdirectories on the appliance.
  - **Allow file transfers from:** Configure the appliance to grant or deny permissions to FTP file distribution requests from clients, as described in *Configuring Access Control for File Distribution* below.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

### Enabling FTP Anonymous User

The 'anonymous' FTP login is disabled by default, except when upgrading an earlier version in which case anonymous FTP is automatically enabled.

When you enable anonymous FTP at Grid level, you enable anonymous FTP on all Grid members running the FTP service. Anonymous user is only allowed to download files, even if the member is enabled to allow uploads.

1. From the **Data Management** tab, select Grid **File Distribution** Properties on the toolbar.
2. In the *Grid File Distribution Properties* dialog box, select the **Enable Anonymous FTP** checkbox.
3. Click **Save & Close**.

### Configuring the HTTP Service

To allow file distribution access using HTTP, you must specify clients that can request the service and then enable the HTTP service on the appliance.

Before you enable the HTTP service, however, be aware of the following configuration rules:

- HTTP only runs on the active member of an HA pair.
- HTTP can run on the master or any member.
- HTTP always runs on the LAN port, never the MGMT port.
- HTTP to HTTPS redirect becomes non-functional if the file distribution service is enabled and all administrative access is run on the LAN port. For more information on HTTP redirect, see [Enabling HTTP Redirection](#). For information on how to specify the MGMT port for HTTP, see [Using the MGMT Port](#). To configure the HTTP file distribution service on a member:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.
2. In the *Member File Distribution Properties* editor, select the **HTTP** tab, and then complete the following:
  - **Allow Any:** This is selected by default to allow HTTP file distribution requests from any client.
  - **Only these addresses:** Select this to configure the access control list for allowing HTTP file distribution requests from clients, as described in *Configuring Access Control for File Distribution* below.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

### Configuring Access Control for File Distribution

You can select a named access control list (ACL) or create individual access control entries (ACEs) for each file distribution service (TFTP, FTP, HTTP) to control access to file distribution requests from specific clients. You can grant or deny access from specific IPv4 addresses and IPv4 networks, but you cannot do so for IPv6 addresses and IPv6 networks as well as TSIG key based ACEs.





#### Note

For HTTP service, you can grant permissions to all clients or specific clients, but you can deny permissions only to all clients, not specific clients.

When you grant access to a network for a specific file distribution service, all clients in the network are allowed to request file distribution service. You can deny services to specific IP addresses within the network by adding these addresses to an access control list and denying access to the service. Ensure that you list these IP addresses before the network address in the list because the appliance applies permissions to the addresses in the order they are listed. You can use the arrow keys to move the addresses up and down the list after you add them. For information about how to create a named ACL, see [Configuring Access Control](#).

To configure an access control list for a file distribution service:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.
2. In the *Member File Distribution Properties* editor, select a service tab: **TFTP**, **FTP**, or **HTTP**.
3. In the **Allow these clients to perform file transfers** section, select one of the following:
  - **Any**: Select this to allow any clients to use the HTTP file distribution service. This is selected by default.
  - **Named ACL**: Select this and click **Select Named ACL** to select a named ACL that contains only IPv4 addresses and networks. File distribution does not support IPv6 addresses/networks and TSIG key based ACEs. When you select this, the appliance allows clients that have the **Allow** permission in the named ACL to use the file distribution service. You can click **Clear** to remove the selected named ACL.
  - **Set of ACEs**: Select this to configure individual access control entries (ACEs). Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding.
    - **IPv4 Address**: Select this to add an IPv4 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
    - **IPv4 Network**: In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address**: Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
      - **Permission**: Select **Allow** or **Deny** from the drop-down list.
    - **Any Address/Network**: For TFTP and FTP only. Select this to allow or deny access to all IPv4 addresses and networks. The default permission is **Allow**, which means that the appliance allows access to and from all IPv4 clients. You can change this to **Deny** to block access.After you have added access control entries, you can do the following:
    - Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
    - Reorder the list of ACEs using the up and down arrows next to the table.
    - Select an ACE and click the Edit icon to modify the entry.
    - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

#### Modifying Access Control Lists

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.
2. In the *Member File Distribution Properties* editor, select the tab of the service to which the list belongs.
3. In the **Allow file transfers from** section, modify the fields as described in [Configuring Access Control for File Distribution](#) above.

You can also do the following:

- Add a new permission. For information, see [Configuring Access Control for File Distribution](#) above.
- Delete a permission by selecting it and clicking the Delete icon.
- Reorder the list by selecting a permission and clicking an arrow next to the list to move the permission up or down the list.

## Starting and Stopping File Distribution Services

You can enable and disable a file distribution service on a specific Grid member or on multiple members. You must have read/write permission to the Grid members to start and stop a service on them.

Starting a service on a member:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* checkbox, and then click the Start icon from the Toolbar. You can select multiple members by selecting their checkboxes.
2. From the Start drop-down menu, select the service you want to start.
3. In the *Start Service* dialog box, click **Yes**.  
Grid Manager enables the selected service on the selected member and displays the service status in the Status column in the panel.

Stopping a service on a member:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab -> *member* checkbox, and then click the Stop icon from the Toolbar. You can select multiple members by selecting their checkboxes.
2. From the Stop drop-down menu, select the service you want to stop.
3. In the *Stop Service* dialog box, click **Yes**.  
Grid Manager disables the selected service on the selected member and displays the service status in the Status column in the panel.



### Note

When you start or stop a service, there may be a short delay before Grid Manager displays the correct status.

## Monitoring File Distribution Services

To view the current status of the file distribution services:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab.
2. Grid Manager displays the following information:
  - **Name:** The name of the Grid member.
  - **Address:** The IP address of the Grid member.
  - **Status:** The overall status of the file distribution services running on the member. You can mouse over on the field to view the status of each service. This field can display one of the following:
    - **Not Running:** All the file distribution services are disabled.
    - **Running:** One or more of the file distribution services are running properly.
    - **Warning:** The services are functioning properly. However, there are some issues, such as storage space has reached 90%, about the services.
    - **Error:** One or more of the services have service issues.
  - **Comment:** Information about the member.
  - **Site:** The location to which the member belongs. This is one of the pre-defined extensible attributes.

You can sort the information in ascending or descending order by columns. You can also print and export the information in this panel.

## Managing Directories

You can create directories on the Grid Master and on Grid members, in which you can store your files. You can manage the directories in the following ways:

- Create a directory structure for file distribution, as described in [Adding Directories](#) below.
- Modify the directory name and permissions, as described in [Modifying Directories](#) below.
- Create a Virtual TFTP root directory, as described in [Creating a Virtual TFTP Root Directory](#) below.
- View the directories, as described in [Viewing Directories From the Files Tab](#) below.

### Adding Directories

To add a directory:

1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
2. Click the parent directory link, and then click **Add -> Directory** from the Toolbar.
3. Grid Manager adds a new directory to the parent directory and gives it the default name **NewDirectory**. You can modify the directory name and permissions, as described in [Modifying Directories](#) below.

### Modifying Directories

1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
2. Select a directory checkbox and click the Edit icon.
3. The *Directory* editor provides the following tabs from which you can modify data:
  - **General**: You can modify the directory name here, except for the Root directory.
  - **Virtual TFTP Root**: You can add an IP Address, a Network or a Range of IP addresses to support VMware ESX hosts who need different PXE boot images based on where they are in the network.
  - **Permissions**: You can add or delete admin permissions in this tab. For information, see [About Administrative Permissions](#).
4. Save the configuration and click **Restart** if it appears at the top of the screen. You can also select a directory and click the Delete icon to delete it.



#### Note

When you delete a directory, the appliance automatically removes all its contents in that directory.

### Creating a Virtual TFTP Root Directory

This section describes how to create a Virtual TFTP root directory for a specific IP address, network, or range of IP addresses. Note that Virtual TFTP root is supported only for file downloads, but not for file uploads using TFTP client.



#### Note

Root directory can not be a virtual TFTP root.

1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
2. Select the directory checkbox and click the Edit icon.
3. In the *Directory* editor select **Virtual TFTP Root**. Click the Add icon and select one of the following:
  - **IP Address**: This creates a virtual TFTP directory that the clients from a specified IP address will see as the root directory.
  - **Network**: This creates a virtual TFTP directory that the clients on a specified network will see as the root directory.

- **Range:** This creates a virtual TFTP directory that the clients in a specified range of IP addresses will see as the root directory.
4. From the drop-down in the **Member** column, select the member on which to make the virtual TFTP root directory.
  5. In the **Address/Network** column, enter a value:
    - **IP Address:** Enter the IP address of the client that will have access to the virtual TFTP root directory. This IP address must be on the allow list in the TFTP access control list.
    - **Network:** Enter a network address using the format 10.0.0.0/24. This allows all clients in that network to access the virtual TFTP root directory. This network address must be on the allow list in the TFTP access control list.
    - **Range:** Enter the first IP address in the range Address/Network column, and the last IP address in the range in the End column. This allows all clients in that range to access the virtual TFTP root directory. This range must be on the allow list in the TFTP access control list.
  6. Click Save & Close. Click **Restart** if it appears at the top of the screen.
  7. To create more virtual TFTP root directories, repeat Steps 3 through 5.

### Viewing Directories From the Files Tab

1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
2. Grid Manager displays the following information in the Root directory.
  - **Name:** The name of the directory or file.
  - **Type:** Depending on the file type, this can be **Directory** or **File**.
  - **Size:** The file size in B, KB, or MB depending on whether the file size crosses the unit limit or not. For example, if the file size is 1023, Grid Manager displays 1023 B. If the file size is 1025, Grid Manager displays 1 KB. For a directory, Grid Manager displays a dash (-).
  - **Date Modified:** The timestamp when the directory was last created or when the file was last modified.
3. Click the directory link to view files and directories in a specific directory. You can also do the following in this panel:
  - Sort the information in ascending or descending order by columns.
  - Use the breadcrumb to go to a specific directory.
  - Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
  - Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
  - Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).
  - Print and export the information in this panel.
  - Add a directory or a file. For information, see [Adding Directories](#) and [Managing Files](#).
  - Open and edit a directory. For information, see [Modifying Directories](#) above.

## Managing Files

This section describes how to upload files using the Grid Manager or a file transfer client. You can upload files to the Grid Master or to individual members.

### Uploading Files

Some things to keep in mind when you upload files:

- When you use the Grid Manager to upload files, you can upload files only to the Grid Master, not to individual members of the Grid.
- To upload files to a member, you must use an FTP client, TFTP client or HTTP client. Files uploaded by file transfer clients to any member, will be synchronized back to Grid Master.
- The logs for file transfers using third party clients can be found in syslog.
- You can use a third party file transfer client to upload and retrieve files:
  - If the 'anonymous' login is enabled, you can retrieve files but this 'anonymous' user can not upload files even if the "Allow uploads" option is enabled.

- If you create a user to use with a third party transfer client, this must be an FTP user with read/write permissions in their directory.
- You can upload a maximum of 10,000 files.
- If uploading a file exceeds the storage limit of 2 GB for a single file, NIOS logs a message and does not upload the file. For information about file distribution storage, see [Modifying File Distribution Storage Limits](#).
- If you upload a file that has the same name and path as an existing file, NIOS automatically replaces the old file.



#### Note

Administrators with superuser privileges can manage uploading files. Limited-access admins with read/write permissions to specific directories can upload files to the directories. For information, see [Administrative Permissions for File Distribution Services](#).

### Enabling Upload to Grid Members

1. From the **Data Management** tab, select Grid **File Distribution** Properties on the toolbar.
2. In the *Grid File Distribution Properties* dialog box, select the **Allow Upload to Grid Members** checkbox.
3. Click **Save & Close**.

### Uploading Files using Grid Manager

The Grid Manager uploads files only to the Grid Master. The Grid Master propagates the files to the members. You must use a third party file transfer client to upload files directly to an individual member:

- If the 'anonymous' login is enabled, you can retrieve files but this 'anonymous' user can not upload files even if the "Allow uploads" option is enabled.
  - If you create a user to use with a third party transfer client, this must be an FTP user with read/write permissions in their directory.
1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
  2. Select the destination directory link.
  3. Click the **Add** icon -> **File** from the Toolbar.
  4. Select the **Extract files after upload (.zip, .tar, .gz, .tgz)** checkbox in the *Upload* dialog box if you are uploading .zip, .tar, .gz, or.tgz files and you want to automatically extract the files upon upload. Note that the directory structure in the compressed file is restored when the files are extracted. A directory that already exists it will be replaced by an extracted directory with the same name.
  5. Click **Select to** navigate to the file you want to upload.
  6. Select the file you want to upload, and then click **Open**.
  7. If you want to upload more than one file, repeat Steps 4 and 6 until you have selected all the files you want to upload. You can upload a maximum of ten files at one time. Note you can delete an incorrect file selection by clicking the red icon next to the filename before you click Upload.
  8. To verify the upload was successful. roll the mouse cursor over the green check mark next to the file name. If the upload was successful, the message "Upload succeeded." appears.

### Uploading Files Using TFTP, FTP, or HTTP File Transfer Client

You can upload files to the Grid Master or to individual members using a third party FTP client. Files uploaded by file transfer clients to any member, will be synchronized back to Grid Master.

To upload files to a member, you must first enable the **Allow Upload to Grid Members** checkbox in the *Grid File Distribution Properties* dialog box. See [Enabling Upload to Grid Members](#).

You must add an FTP user before you can upload files using a third party FTP client. This must be an FTP user. It is not the NIOS admin. For information see [Adding FTP Users through Grid Manager](#).

## Deleting Files From the Grid Master

If the FTP user has read/write permissions, then that user can delete files from the Grid member wherever that FTP user is connected. Only files can be deleted but not directories.

1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
2. Grid Manager displays the files and folders in the Root directory. Click the directory link to see the files in a specific directory.
3. To delete a file, select the checkbox and then click the **Delete** icon.

## Deleting Files From a Member

You can delete files from a member only if "**No**" appears in the *Synchronized with Grid Master* column.

If the FTP user has Read/Write permissions, then that user can delete files from the Grid member wherever that FTP user is connected. Only files can be deleted but not directories.

1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
2. Grid Manager displays the files and folders in the Root directory. Click the directory link to see the files in a specific directory.
3. If "**No**" appears in the *Synchronized with Grid Master* column, select the checkbox, then click the **Delete** icon.

## Viewing Files

You can view files from the Files Tab and from the Members Tab.

### Viewing Files from the Files Tab

1. From the **Data Management** tab, select the **File Distribution** tab -> **Files** tab.
2. Grid Manager displays the following information in the Root directory.
  - **Name:** The name of the file.
  - **Type:** Depending on the file type, this can be **Directory** or **File**.
  - **Size:** The file size in B, KB, or MB depending on whether the file size crosses the unit limit or not. For example, if the file size is 1023, Grid Manager displays 1023 B. If the file size is 1025, Grid Manager displays 1 KB. For a directory, Grid Manager displays a dash (-).
  - **Date Modified:** The timestamp when the directory was last created or when the file was last modified.

You can view files and directories in a specific directory by clicking the directory link.

### Viewing Files from the Members Tab

1. From the **Data Management** tab, select the **File Distribution** tab -> **Members** tab.
2. Grid Manager displays the following information in the Root directory.
  - **Name:** Member name.
  - **IPv4Address:** Member's IP address.
  - **Status:** State of the member, running or not running.
  - **Comment:** Additional comments about the member.
  - **Site:** User defined information about the site.
3. To see the files on the Grid Master, click on the name of the Grid Master. Grid Manager displays the following information:
  - **Name:** The name of the file.
  - **Type:** Depending on the file type, this can be **Directory** or **File**.
  - **Size:** The file size in B, KB, or MB depending on whether the file size crosses the unit limit or not. For example, if the file size is 1023, Grid Manager displays 1023 B. If the file size is 1025, Grid Manager displays 1 KB. For a directory, Grid Manager displays a dash (-).
  - **Date Modified:** The timestamp when the directory was last created or when the file was last modified.

**Tip:** When you drill down on the Grid Master from the Members tab, the Add icon is activated.

4. To see the files on a member, click on the name of the member. Grid Manager displays the following information:
  - **Name:** The name of the file.
  - **Type:** Depending on the file type, this can be **Directory** or **File**.
  - **Size:** The file size in B, KB, or MB depending on whether the file size crosses the unit limit or not. For example, if the file size is 1023, Grid Manager displays 1023 B. If the file size is 1025, Grid Manager displays 1 KB. For a directory, Grid Manager displays a dash (-).
  - **Date Modified:** The timestamp when the directory was last created or when the file was last modified.
  - **Synced with Grid Master:** You cannot delete files with a value other than "No". If this value is "No", you must delete the file.



**Note**

You cannot upload, modify, or delete a file or a directory when you drill down from the Members tab.

You can also do the following in this panel:

- Sort the information in ascending or descending order by columns.
- Use the breadcrumb to go to a specific directory.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).
- Print and export the information in this panel.
- Add a directory or a file. For information, see [Adding Directories](#) and [Managing Files](#).
- Open and edit a directory. For information, see [Modifying Directories](#).
- You can enable "group by" function to organize your members in a meaningful way and quickly identify them based on common data. Complete the following to group members with the same extensible attribute value:
  - **Group Results:** Select this checkbox to enable the appliance to group members by extensible attributes.
  - **Group By:** From the drop-down list, select the first extensible attribute that you want the appliance to use for filtering members.  
Grid Manager displays data per group of members configured with the same extensible attribute value. For more information about grouping by extensible attributes, see [Grouping Members by Extensible Attributes](#).

## Managing Users

This section describes how to add and modify user accounts for use with an FTP client.

You must be a NIOS admin user with super user privileges to add, modify, or delete FTP users. FTP users are created at Grid level, so the same users will be available to access FTP service on all members.

- Each user must have unique Username.
- By default, the home directory with the user name is under the `/ftpusers` directory. However, the user can also choose to use an existing directory outside of `/ftpusers` as his home directory. If the admin specified home directory is not available then it will raise an error.
- Permission: Read-write or Read-only are assigned for each FTP user. Users with read-write permissions are allowed to upload files, delete files and list the files and directories under his home directory.
- You can have multiple users to use same home directory. One user may have read-only permissions while others have read-write permissions on same home directory.
- FTP users are not allowed to add, modify, or delete the directories, even with read-write permissions.

- If the "Allow uploads on the member" is disabled, then users with read-write permission also can not upload files to his home directory.

## Users Default Home Directory

- FTP users default home directory is `/ftpusers`.
- The `/ftpusers` directory is created by default and listed in the 'Files' viewer under the root directory. By default, home directories for FTP users are under this directory.
- NIOS admin is allowed to upload and delete files to and from users home directories.
- Files uploaded by FTP users are visible in the Grid Manager.

## Adding FTP Users through Grid Manager

1. From the **Data Management** tab, select the **File Distribution** tab -> FTP Users tab, and then click the Add icon.
2. In the *Add FTP User* dialog box, complete the following:
  - **Username:** Enter a name for the user. This is the username that the user uses to log in.
  - **Password:** Enter a password for the user to use when logging in.
  - **Confirm Password:** Enter the same password.
  - **Permissions:** From the drop down choose from the following:
    - **Read Only:** This allows the user to display files, but not to upload, or delete files using a third party FTP client.
    - **Read/Write:** This allows the user to upload, delete and list files using a third party FTP client.
  - Choose a directory for this user. This is the directory where files uploaded with this Username will go:
    - **Create Home Directory:** This creates a directory using the Username.
    - **Choose Specified Directory:** This allows the user to choose an existing directory.
3. Click **Save & Close**.

## Adding FTP Users through CSV Import

You can add an FTP User by importing a CSV file with the headers in the following format:

```
version 1.0,,,,
header-ftpuser,username*,password*,create_home_dir,home_dir,permission
ftpuser,user1,passwd1,True,/ftpusers/user1,R0
```

## Modifying FTP Users

1. From the **Data Management** tab, select the **File Distribution** tab -> FTP Users tab.
2. Select the checkbox for the user you want to modify and click the Edit icon.
3. In the *FTP User Editor* you can modify the following:
  - **Password:** Enter a password for the user to use when logging in.
  - **Confirm Password:** Enter the same password.
  - **Permissions:** From the drop down choose from the following:
    - **Read Only:** This allows the user to display the files and their properties, but not to edit them.
    - **Read/Write:** This allows the user to display and edit the files and their properties.
4. Click **Save & Close**.

## bloxTools Environment

The bloxTools environment provides a pre-installed environment for hosting custom web-based applications. This section includes the following topics:

- [About the bloxTools Environment](#)
- [Using the bloxTools Environment](#)



- [Monitoring the Service](#)

## About the bloxTools Environment

The bloxTools environment provides tools for creating custom applications that facilitate the administrative tasks in your organization. It provides a pre-installed environment for running applications using Perl, Python, PHP, CGI scripting, and Infoblox API libraries. Note that no direct external remote user (telnet and ssh, for example) or shell access is available in this environment.

The bloxTools environment "borrows" resources such as CPU, memory, disk space, and networking from the host Infoblox appliance, but is logically separated from the NIOS. The logical separation ensures that any failure in the bloxTools service does not affect the other services running on the appliance.

The bloxTools environment can only be configured to run on an independent appliance or a Grid member. You cannot run the bloxTools service on a Grid Master or a Grid Master candidate.



### Note

In previous NIOS releases, you could run the bloxTools service only on a Grid Master. If bloxTools has been configured to run on a Grid Master before an upgrade, the bloxTools service continues to run on the Grid Master after an upgrade. This configuration is preserved mainly for migration purposes only. Infoblox strongly recommends that you move the bloxTools service to a Grid member after the upgrade. For information about moving the bloxTools service, see [Using the bloxTools Environment](#).

In a Grid, you can run the bloxTools service only on one Grid member at a time, and you cannot configure this member as a Grid Master candidate. However, you can move the bloxTools service from one member to another. For information about moving the bloxTools service, see [Using the bloxTools Environment](#).

On an HA member, the bloxTools service runs on the active node. If there is an HA failover, the bloxTools service is automatically launched after the passive node becomes active. For information, see [About HA Pairs](#).



### Note

When you run the bloxTools service on an independent appliance or a Grid member, the performance of other services running on the appliance may be affected. Infoblox recommends that you run the bloxTools environment on a member that does not host critical services.

After you enable the bloxTools service and configure its built-in file transfer services, you can upload content to the bloxTools portal using either an FTP (File Transfer Protocol) or SFTP (SSH File Transfer Protocol) client. The uploaded content is included in system backups and you can restore it from the backups.

If you have further questions about bloxTools, visit the community site at <https://community.infoblox.com>.

## System Requirements

The following table shows which Infoblox physical appliances support the bloxTools service.

Supported Infoblox Appliance
TE-815
TE-1415
TE-2215

The following table shows which Infoblox appliances support the bloxTools service and the memory requirement for each. The service "borrows" host resources such as CPU, memory, and disk space from the host Infoblox appliance.

Supported Infoblox Appliance	Memory Requirement
IB-V1410 IB-V1415 IB-V1420 IB-V1425	128 MB to 2048 MB The default is 256 MB
IB-V2210 IB-V2215 IB-V2220 IB-V2225	128 MB to 2048 MB The default is 256 MB
IB-V4015 IB-V4025	128 MB to 4096 MB The default is 256 MB

## Using the bloxTools Environment

Complete the following tasks to upload custom applications to the bloxTools environment:

1. Log in to the appliance as a superuser and configure the bloxTools service, as described in [Configuring the Service](#) below.
2. Use an FTP or a SFTP client to upload content to the bloxTools environment.

In addition, you can schedule tasks as described in [Scheduling Tasks](#) below, and monitor the bloxTools service as described in [Monitoring the Service](#).

---

**WARNING:** *Resetting the Grid member using either the `reset all` or `reset database` CLI commands permanently deletes the content you uploaded to the bloxTools environment. Infoblox recommends that you backup the appliance before using any of these commands.*

---

## Configuring the Service

When you configure the bloxTools service, you can enable FTP, SFTP, and HTTPS, and set their operational parameters. FTP and SFTP are the services you use to upload data. You can disable these services when they are not in use. HTTPS must remain enabled to allow the web based bloxTools applications to run. Note that the bloxTools service uses the same SSL certificate as the host Infoblox appliance. For information on certificates, see [Managing Certificates](#). You can configure the bloxTools service on port 443, 444 or on a port between 1024 to 63999. You can also enable HTTP to HTTPS redirection for the bloxTools service from the default HTTP port 80 to any specified HTTPS port. When you enable HTTP to HTTPS redirection, all the requests sent to the HTTP port are redirected to the HTTPS port configured for the bloxTools service. By default, NIOS appliance does not redirect HTTP requests to HTTPS.



### Note

When you redirect HTTP to HTTPS, the connection is not as secure as compared to connecting directly through HTTPS.

Note the following when you configure the bloxTools service on port 443:

- HTTP to HTTPS redirection from Grid member to Grid Master is disabled.
- HTTP file distribution service is not allowed.

In previous NIOS releases, you could run the bloxTools service on a Grid Master or a Grid Master candidate. If you have not removed the bloxTools service from a member before you upgrade it to NIOS 6.11.0 or later, you cannot configure the bloxTools service on port 443 until you move the bloxTools service to an upgraded Grid member. For information, see [Moving the bloxTools Service](#).

To configure the bloxTools service:

1. Log in as a superuser.
2. From the **Grid** tab, select the **Grid Manager** tab, and then click **bloxTools**. In the **Services** tab, click **Edit** -> **Grid bloxTools Properties** from the Toolbar.
3. In the *Grid bloxTools Properties* editor, complete the following:
  - **Enable Web Service:** Select **HTTPS Port** to enable users to access the applications through an HTTPS connection. The default port is 444. You can change the default HTTPS port to 443 or to a port between 1024 to 63999.
  - **Redirect Bloxtools HTTP to HTTPS:** Select this checkbox to enable redirection from the default HTTP port to the HTTPS port. The default HTTP port is 80. This is disabled by default.
  - **Enable FTP Service:** Select **FTP Port** to enable the FTP service. The default port is 26. You can change the port number to suit your environment.
  - **Enable SFTP Service:** Select **SFTP Port** to enable the SFTP service for secure file transfer. The default port is 28. You can change the port to a number between 1024 and 63999, provided that the port is not currently used for another purpose.
  - **Login:** Enter the username for the FTP and SFTP services. The username can contain lower case letters, numbers, underscores (\_), and dollar signs (\$), and it must begin with a letter, not a number.
  - **Set Password:** Enter the password for the FTP and SFTP services in this field.
  - **Retype Password:** Enter the same password.

Note that the password is sent as clear text when you use the FTP service. To maintain security on the Infoblox appliance, this password should be different from the password set for the Infoblox appliance.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

When you configure the bloxTools service on an independent appliance, you can configure the allocated memory in the *System bloxTools Properties* editor. For information, see [Allocating Memory](#) below.

## Allocating Memory

You can configure the memory you want to allocate to the bloxTools service. You must configure this at the member level. If you run the bloxTools service on an independent appliance, you can configure the allocated memory in the *System bloxTools Properties* editor.

To configure the allocated memory:

1. Log in as a superuser.
2. From the **Grid** tab, select the **GridManager** tab, and then click **bloxTools**. In the **Services** tab, click **Edit** -> **MemberbloxToolsProperties** from the Toolbar.
3. In the *MemberbloxToolsProperties* editor, complete the following:
  - **Allocated Memory (MB):** The service "borrows" host resources such as CPU, memory, and disk space from the host Infoblox appliance. The default amount of memory the appliance allocates for the bloxTools environment is 256 MB. You can change this allocation, depending on the appliance platform. See [System Requirements](#) for the requirements and allowed values of each appliance.

## Uploading Files

Use an FTP or a SFTP client to upload content, such as Perl modules, JavaScript files, PHP files, CGI files, and image files, to the bloxTools environment. You can upload a maximum of 4 GB of data. After you have uploaded content to your bloxTools environment, you should disable the FTP and SFTP services to prevent unauthorized or accidental changes. To upload files using the FTP service:

1. Open an Internet browser window and log in to the FTP service by entering:

```
ftp://Grid_member_ip_addr:ftp_port
```

For example, if the IP address of the Grid member is 10.1.1.1 and the FTP port number is 26, enter: `ftp://10.1.1.1:26`

2. In the *Authentication Required* dialog box, enter the username and password. This is the username and password you entered for the FTP service in the *bloxTools Environment* editor on the appliance.
3. Follow the instructions provided by your FTP client to upload the files.

To upload files using the SFTP service:

1. Open a terminal window and log in to the SFTP service by entering:

```
sftp -oPort=sftp_port sftp_user@Grid_member_ip_addr
```

For example, if the IP address of the Grid member is 10.1.1.1, the login username for the SFTP service is jdoe, and the SFTP port number is 28, enter:

```
sftp -oPort=28 jdoe@10.1.1.1
```

2. Enter the password. This is the password you entered for the SFTP service in the *bloxTools Environment* editor on the appliance.
3. Follow the instructions provided by your SFTP client to upload the files.



#### Note

On a computer running Microsoft Windows, you can use WinSCP as the FTP or SFTP client for uploading files. The *bloxTools* environment stores the uploaded data in the `/portal` directory.

## Scheduling Tasks

*bloxTools* includes support for the Perl module `Config::Crontab` so you can manage scheduler services. You can use the scheduler to execute commands in the future. You can also schedule recurring commands. For example, you can schedule the creation of a host record or schedule recurring reports. The scheduler allows default "user level" crontab access and you can use the user account 'nobody' to submit commands. The Grid Master replicates the crontab data to the Master Candidates.

## Moving the *bloxTools* Service

In a Grid, you can move the *bloxTools* service from one Grid member to another. When you move the *bloxTools* service, the source member synchronizes data with the Grid Master, and the Grid Master synchronizes data with the destination member. The time to resynchronize the *bloxTools* data on to the destination member depends on the amount of data to synchronize and the Grid configuration. If the migration takes longer than two minutes, it becomes a long running task. This allows the move of the *bloxTools* service to run in the background while you perform other tasks. Note that on an independent appliance, you cannot move the *bloxTools* service to another member.

After an upgrade from previous NIOS releases, Grid Manager displays a warning message in the system message panel if you have previously configured to run the *bloxTools* service on the Grid Master. You can click **Move** in this panel to launch the *Move bloxTools* dialog box to move the *bloxTools* service to a Grid member.

To move the *bloxTools* Service:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **bloxTools** -> **Move** from the Toolbar.
2. In the *Move bloxTools* dialog box, complete the following:
  - **Source Member**: Displays the name of the Grid member that is currently running the *bloxTools* service. You cannot modify this field.
  - **Destination Member**: Click **Select**. In the *Member Selector* dialog box, select the member to which you want to move the *bloxTools* service. Grid Manager displays the name of the selected member here.
3. Click **Move**.

The appliance synchronizes data with the Grid Master, and the Grid Master synchronizes data with the destination member. This may take a while to complete depending on your Grid configuration and the amount of data.

## Monitoring the Service

Infoblox provides several tools for monitoring the bloxTools Environment. The bloxTools Environment has its own syslog service which you can access to view logs generated by the bloxTools service and its processes. The *Detailed Status* panel also displays the status of the bloxTools Environment.

### Viewing the Logs

The bloxTools Environment generates the following logs:

- `access.log`: The Apache access log
- `error.log`: The Apache error log
- `syslog.log`: The bloxTools Environment system log

These log files are included in the support bundle. You can download the log files using FTP. You can also connect to the CLI of the member running the bloxTools environment and use the following commands to view the logs:

- Use the `show file` command to view the list of log files.
- Use the `show bloxtools` command to view the status of the bloxTools Environment.
- Use the `show file bloxtools portal_access` command to view the web portal access log.
- Use the `show file bloxtools portal_error` command to view the web portal error log.
- Use the `show file bloxtools portal_log` command to view the web portal system log.





### Viewing Detailed Status

You can view the status of the bloxTools Environment from the **Services** tab of the **Grid Manager** tab. To display the bloxTools service status, from the **Grid** tab, select the **Grid Manager** tab -> **Services** tab, and then click **bloxTools**. Grid Manager displays all Grid members that can host the bloxTools service. The name of the Grid Master is displayed only if you have completed an upgrade and previously configured the bloxTools service to run on the Grid Master. Though you can continue to run the bloxTools service on the Grid Master, Infoblox strongly recommends that you move the bloxTools service to a Grid member. For information, see [Moving the bloxTools Service](#).

Grid Manager displays the following information about all Grid members:

- **Name**: The Grid member name.
- **Service Status**: Indicates the current operational status of the bloxTools service running on the member. This can include the migration status if you are moving the bloxTools service to another member.
- **IP Address**: The IP address of the member.
- **Comment**: Information about the bloxTools Environment.
- **Sit**: The location to which the member belongs. This is one of the predefined extensible attributes.

The service status icon indicates the operational status of the bloxTools Environment and the usage percentages for the CPU, memory and disk resources. The status icon can be one of the following:

Icon	Color	Meaning
	Gray	The bloxTools Environment is disabled or offline.
	Green	The CPU, memory, and disk usage is below 80%.
	Yellow	Usage of at least one of the following resources is greater than or equal to 80%: CPU, memory or disk. The description indicates the percentage of each resource.
	Red	The bloxTools Environment is down, or an essential service within the bloxTools Environment has failed.

## RIR Registration Updates

This section explains how to configure the Infoblox Grid to manage RIR (Regional Internet Registries) allocated addresses and submit registration updates to the RIPE database. It includes the following topics:

- [RIR Address Allocation and Registration Updates](#)
- [Configuring RIR Registration Updates](#)
- [Managing RIR Data](#)
- [Managing RIR Attributes](#)
- [Monitoring RIR Data](#)
- [Requirements and Permissions](#)

---

### RIR Address Allocation and Registration Updates

You can configure the Infoblox Grid to manage allocated IP address blocks that ISPs (Internet Service Providers) receive from their RIRs (Regional Internet Registries). An RIR is an entity that manages the Internet number resources, which include IP addresses and autonomous system numbers, within a specific region of the world. RIRs use SWIP (Share WHOIS Project) or RWhois (Referral WHOIS) servers to provide address allocation information for IP address blocks. Typically, an RIR determines the address blocks to be allocated for specific organizations (typically ISPs), while an ISP manages the allocated address blocks, associated organizations and corresponding RIR registrations. An organization can determine when to request for more address blocks from its RIR. Most ISPs manage multiple organizations and synchronize network address data with their RIRs every few months.

To leverage IPAM (IP Address Management) on the NIOS appliance, you can enable the Infoblox Grid to manage RIR allocated addresses and send registration updates to the RIPE (Réseaux IP Européens) database as often as you update RIR data on NIOS. RIPE is one of the five RIRs in the world that manages the allocation and registration of Internet number resources for Europe, Russia, the Middle East, and Central Asia.



#### Note

The RIR registration update feature is not supported in a Multi-Grid configuration.

### About the RIPE Database

The RIPE database contains registration details of IP addresses and AS numbers originally allocated by the RIPE NCC (RIPE Network Coordination Center). The database contains information such as organizations that hold IP resources, where the allocations were made, and contact details for the networks. Organizations or individuals that hold the allocated address blocks are responsible for updating information in the database.

The NIOS appliance supports submitting registration and reassignment updates to the RIPE database, which can be accessed through the RIPE API interface or an email template. For more information, see *Configuring RIR Communication Settings*.



#### Note

Before the NIOS appliance sends registration updates to the RIPE database, it does not validate the data you submit. Therefore, if you enter invalid information that cannot be mapped to the RIPE database, your updates will fail. In addition, the NIOS appliance does not synchronize data from the RIPE database

## Requirements and Permissions

To manage RIR allocated addresses, organizations, and network utilization that contain RIR assignments, you must first enable support for RIR registration updates, and then configure the RIR communication method. Note that once you have enabled support for RIR registration updates, settings and fields that are relevant to this feature are enabled in Grid Manager. You do not need a special license to use this feature.

Only superusers can create, modify, and delete RIR organizations. Limited-access users can manage RIR allocated address blocks if they have the required permissions to the objects.

To view and manage RIR related data, admins must have permissions to the applicable resources. For example, to view RIR networks, admins must have read-only permission to the networks; and to edit them, admins must have read/write permission to them. For more information about admin permissions, see [About Administrative Permissions](#).

## Configuring RIR Registration Updates

To manage RIR allocated addresses and send registration updates through NIOS, you first add RIR organizations and create RIR allocated networks in NIOS. You can then reassign network addresses within the RIR allocated address block to other organizations based on your requirements, and then configure NIOS to send registration updates directly to the RIPE database. Any data you manage through the Grid is handled by the Grid Master.

When the Grid Master is an HA pair, the active node handles the submission of data. If an HA failover occurs during a submission, the failing node immediately aborts the submission. The new active node resumes the next submission. For information about HA pairs, see [About HA Pairs](#).

To manage and submit updates to the RIPE database, you must first enable the Grid to support RIR registration updates. You can then enter RIR information, such as RIR organizations and RIR attributes.

To configure the Grid to manage RIR allocated addresses and submit updates to RIPE, complete the following:

1. Enable support for RIR registration updates, as described in [Enabling Support for RIR Registration Updates](#) below .
2. Define the method to communicate updates to RIPE, as described in [Configuring RIR Communication Settings](#) below.
3. Add and configure RIR organizations and RIR organizational attributes, as described in [Adding RIR Organizations](#), see [Managing RIR Data](#).
4. Add allocated address blocks and assign specific network addresses to RIR organizations, as described in [Adding and Assigning RIR Networks](#), see [Managing RIR Data](#).
5. Review and submit registration updates to RIPE, as described in [Previewing Registration Updates](#), see [Managing RIR Data](#). You can also perform the following tasks:
  - View a list of RIR organizations, as described in [Viewing RIR Organizations](#).
  - Modify RIR organizations and RIR organizational attributes, as described in [Modifying RIR Organizations](#).
  - Monitor network utilization for networks that contain RIR assignments. For information about IPAM Home, see [Managing IPv4 Networks](#) or for more information about [Viewing Networks](#), see [Configuring IPv4 Networks](#).

## Enabling Support for RIR Registration Updates

Before you can manage RIR data through Grid Manager, you must first enable support for RIR registration updates. To enable support for RIR registration updates:

1. From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties* editor, select the **General** tab -> **Advanced** tab, and complete the following:
  - **Enable Updates of RIR Registrations**: Select this to enable the support for submitting RIR registration updates to the RIPE database. When you enable this feature, you can configure the appliance to send registration updates to RIPE for network reassignments and reallocations.



Note: Ensure that you configure DNS resolvers for the Grid when you enable this feature. For information about how to configure DNS resolvers, see [Enabling DNS Resolution](#).

3. Save the configuration.

## Configuring RIR Communication Settings

You can configure the appliance to send RIR address updates to RIPE through the RIPE REST API or through an email using the maintainer email address specified in the RIR organization. Note that when you use the API method to delete a registered address block, you do not need to submit RIR attributes that match the data in the RIPE database. However, when you use the email method, you must enter RIR attribute values that match the data in the database. Otherwise, your submission will fail. To view examples of registration updates that NIOS sends, see [Previewing Registration Updates in Managing RIR Data](#).

To configure the RIR communication settings:

1. From the **Administration** tab, select the **RIR** tab, and then select **RIR Settings** -> **RIPE** from the Toolbar.
2. In the *RIR Communication Settings - RIPE* editor, select one of the following to determine how the appliance sends updates to RIPE. The default is **API**.
  - **API**: The appliance sends RIR updates to RIPE through the RIPE API. The default destination is <https://rest.db.ripe.net> for accessing the production database and <https://rest-test.db.ripe.net> for accessing the test database. Click **Override** and enter a different URL to override the default value. When you select this as the communication method, the registration status will be updated automatically after the registration update is completed. Note that RIPE supports only secure connections using HTTPS.
  - **Email**: The appliance sends RIR updates to RIPE through the email address displayed in the field. The default is [auto-dbm@ripe.net](mailto:auto-dbm@ripe.net). Click **Override** and enter a different email address to override the default value. The appliance uses a special email template that includes values of certain RIR attributes. If any of the RIR attribute values do not match the database in the RIPE database, your submission will fail. When you select **Email** as the communication method, ensure that you enable email notifications at the Grid level. For information how to enable email notifications, see [Setting SNMP and Email Notifications in Configuring SNMP](#). Note that when you select this as the communication method, the registration status will not be automatically updated. You can manually change the status. For information, see [Modifying RIR Network Data, in Managing RIR Data](#).
  - **None**: The appliance does not send RIR updates to RIPE.
3. Save the configuration.

## Managing RIR Data

An RIR organization provides information about an entity that has registered a network resource in the RIPE Database. This entity can be a company (such as an ISP), a nonprofit group, or an individual. You can add RIR organizations defined in the RIPE database and start managing their data through NIOS.

After you have enabled support for RIR updates and configure the desired communication method for the updates, you can do the following to manage RIR data:

- Add RIR organizations and their associated data, as described in [Adding RIR Organizations](#) below.
- Add the RIR allocated addresses to NIOS and assign specific address blocks to ISP organizations, as described in [Adding and Assigning RIR Networks](#) below.
- View a list of organization objects, as described in [Viewing RIR Organizations](#) below.
- Review the reassignment information before sending the updates to RIPE, as described in [Previewing Registration Updates](#) below.
- Modify RIR organizational data and attributes, as described in [Modifying RIR Organizations](#) below.
- Modify RIR network data and attributes, as described in [Modifying RIR Network Data](#) below.
- Delete RIR organizations, as described in [Deleting RIR Organizations](#) below.
- Delete delegated addresses from an organization, as described in [Deleting RIR Networks](#) below.



## Adding RIR Organizations

Before you can submit any RIR updates to the RIPE database, you must first add the RIR organization and its corresponding data to NIOS. You can also create additional organizations for ISP customers.

To add an organization:

1. From the **Administration** tab, select the **RIR** tab, and then click **Add -> RIPE Organization**.
2. In the *Add RIPE Organization* wizard, complete the following:
  - **Internet Registry:** The default is **RIPE**. This is the RIR that allocates address blocks to your organization. You cannot change this.
  - **Organization Name:** Enter the name of the organization that holds the resources allocated by RIPE NCC. You can enter up to 256 characters. Enter the name in this format: A list of words separated by white space. A word can be made up of letters, digits, the character underscore "\_", and the character hyphen "-". The first character of a word must be a letter or digit and the last character of a word must be a letter, digit or a period. For example, you can enter  
`SPRINT REGION2.`
  - **Organization ID:** Enter the handle or ID of the organization. You can enter up to 23 characters. Enter the ID in this format: Start with **ORG-** followed by two to four characters, then followed by up to five digits and a source specification. Note that the first digit cannot be "0". The source specification starts with "-" followed by the source name that contains up to nine characters in length. For example, you can enter  
`ORG-CA1-RIPE` or  
`ORG-CB2-TEST`
  - **Maintainer:** Enter the name of the maintainer for this organization. This is required. You can enter up to 256 characters; however, note that the RIPE database has an 80 characters limit for this field. A maintainer is any registrant or person to whom the authority to update has been delegated by another registrant either directly or indirectly, and who holds an identifier that allows updates to be authenticated and authorized. Data entered here must match exactly how the maintainer appears in RIPE. Enter the maintainer name in this format: Use letters, digits, the character underscore "\_", and the character hyphen "-". The first character must be a letter, and the last character must be a letter or a digit. You cannot use the following words (they are reserved by RPSL): any, as-any, rs-any, peer, as, and, or, not, atomic, from, to, at, action, accept, announce, except, refine, networks, into, inbound, outbound. Also note the following: Names starting with certain prefixes are reserved for certain object types. For example, names starting with "as" are reserved for as set names. Names starting with "rs-" are reserved for route set names. Names starting with "rtrs-" are reserved for router set names. Names starting with "fltr-" are reserved for filter set names. Names starting with "prng-" are reserved for peering set names. Names starting with "irt-" are reserved for irt names.
  - **Password:** Enter the maintainer password. This is required. You can enter up to 256 characters.
  - **Retype Password:** Enter the same password.
  - **Maintainer Email:** Enter the originating or source email address of the maintainer. This is required.
3. Save the configuration. Note that you cannot schedule the creation, modification, or deletion of an RIR organization.

**RIR Organizational Attributes:** This table lists all predefined RIR attributes associated with the RIR organization. Click the **Value** field of an attribute in the table to enter a value. The **Required** field indicates whether a value for the corresponding attribute is required.

You can add custom attributes by clicking the Add icon and select an attribute from the drop-down list. You can also delete an RIR attribute by selecting its checkbox and clicking the Delete icon.

For information about the RIR Organizational Attributes table, attributes and how to enter their values, see [Managing RIR Attributes](#).

Note that you cannot leave an optional RIR attribute value empty. If you do not have a value for an RIR attribute, you must delete it from the table. You can enter up to 256 characters for all RIR attributes.

## Modifying RIR Organizations

To modify an RIR organization:

1. From the **Administration** tab, select the **RIR** tab -> *rir\_organization* checkbox, and click the Edit icon.
2. In the *Organization* editor, modify the organization information, as described in Adding RIR Organizations. You can also reorder the list of RIR organizational attributes using the up and down arrows.
3. Save the configuration.

## Deleting RIR Organizations

You can delete an RIR organization that does not have any networks assigned to it. When you delete an RIR organization, the appliance moves it to the Recycle Bin, if enabled. You can later restore the network if needed. For information about the Recycle Bin, see [Using the Recycle Bin](#).

To remove an RIR organization:

1. From the **Administration** tab, select the **RIR** tab -> *rir\_organization* checkbox, and then click the Delete icon.
2. In the *Delete Confirmation (RIR Organization)* dialog box, click **Yes**.

## Adding and Assigning RIR Networks

Before you can assign network addresses within an RIR allocated address block to an organization, you must first add the allocated address block to NIOS. Infoblox supports IPv4 and IPv6 network containers and networks. You can also create network templates that are specific for RIR networks. For information about creating network templates, see [About IPv4 Network Templates](#) and [About IPv6 Network Templates](#).

Note that when you add network containers or networks to NIOS, the appliance does not validate whether the corresponding networks actually exist in the RIPE database. Even though you can create the networks in NIOS, the submission of updates for the network may fail. For example, if you create a child network and the parent network is not registered in RIPE, the registration update will fail.

In addition, each network can only be associated with an RIR in one network view. If you have a network address block registered with RIPE in a specific network view, you must not register the same address block in a different network view. When you enable the support for updates of RIR registrations, Grid Manager displays the appropriate data fields that you can use to add or modify RIR related networks. You can do the following to add IPv4 and IPv6 networks:

- Add RIR allocated IPv4 networks to NIOS, or assign addresses to specific organizations. For information see [Adding IPv4 Networks](#).
- Add RIR allocated IPv6 networks to NIOS, or assign addresses to specific organizations. For information, see [Adding IPv6 Networks](#).
- Add IPv4 network templates that are specific to RIR address allocation. For information, see [About IPv4 Network Templates](#).
- Add IPv6 network templates that are specific to RIR address allocation. For information, see [About IPv6 Network Templates](#).

You can also do the following to modify specific data about the RIR networks:

- Modify RIR allocated or assigned IPv4 networks. For information see [Modifying IPv4 Networks](#).
- Modify RIR allocated or assigned IPv6 networks. For information see [Modifying IPv6 Networks](#).
- Modify IPv4 network templates that are specific to RIR address allocation. For information, see [Modifying IPv4 Network Templates](#).
- Modify IPv6 network templates that are specific to RIR address allocation. For information, see [Modifying IPv6 Network Templates](#).

You can preview the information before the appliance submits updates to the RIPE database. To preview registration updates, click **Preview RIR Submissions** in the *Add IPv4 Network* or *Add IPv6 Network* wizards. For more information, see [Previewing Registration Updates](#).



### Note

You can also add RIR networks through **Task Dashboard**. For information, see [The Tasks Dashboard](#).

After you create an RIR network container or network, you can perform the following:

- Split a network that has an organization ID. A child network that is created does not contain an organization ID by default. You must assign an organization ID to the child network after splitting it. For information about splitting an RIR network, see [Splitting IPv4 Networks into Subnets](#) and [Splitting IPv6 Networks into Subnets](#).
- Resize an IPv4 RIR network that contains an organization ID and has been registered with RIPE. For more information, see [Resizing IPv4 Networks](#).

## Viewing RIR networks

You can view a list of IPv4 and IPv6 RIR networks in the **Data Management** tab -> **IPAM** tab or the **Data Management** tab -> **DHCP** tab -> **Networks** tab -> **Networks** section. For more information, see [IPAM Home](#) and [Viewing Networks](#).

## Modifying RIR Network Data

You can modify certain RIR network information in the **RIR Registration** tab of the *IPv4 and IPv6 Network* editors. To modify RIR network information, complete the following:

1. From the **Data Management** tab, select the **IPAM** tab -> *network* checkbox, and then click the Edit icon.  
or  
From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* checkbox, and then click the Edit icon.
2. In the *IPv4 or IPv6 Network Container* or *Network* editor, click the **RIR Registration** tab, and then complete the following to modify RIR related data for the IPv4 or IPv6 network container or network:
  - **Internet Registry**: Displays the RIR that allocates RIR address blocks. The default is **RIPE**. You cannot change this.
  - **Organization ID**: Displays the organization ID with which this network is associated. You cannot change this.
  - **Registration Status**: Displays the current registration status. This can be **Registered** or **Not Registered**. **Registered** indicates that the network has a corresponding entry in the RIPE database. Note that when you select **API** as the communication method, the registration status will be updated automatically after the registration update is completed. However, when you select **Email** as the communication method, the registration status will not be automatically updated. You can modify this by selecting the appropriate status from the drop-down list.
  - **Status of last update**: Displays the registration status, communication method, timestamp of the last registration update. The status can be Pending, Sent, Succeeded, or Failed. The displayed timestamp reflects the timestamp used on the Grid Master. Each time you send a registration update to create, modify, or delete a network container or network, the updated status and timestamp will be displayed here. If you have selected not to send the registration update, the previous status and timestamp are retained.
  - **Registration Action**: From the drop-down list, select what you want to do with the RIR network updates. If you are creating a top-level network block that has already been assigned to the organization, select **None**. If you are creating a child network within the allocated address block, you can select one of the following:
    - **None**: The appliance does not submit the updates.
    - **Create**: The appliance creates the network container or network for the specified organization.
    - **Modify**: Modifies data for this network container or network.
    - **Delete**: Deletes the RIR network from the organization. When you select this, you must enter a reason for deleting this entry in the **Delete Reason** field.
  - **Do not update registrations**: By default, the appliance sends updates to RIPE if you specify **Create**, **Modify**, or **Delete** as the registration action. Select this if you do not want the appliance to submit updates to the RIPE database.

**RIR Network Attributes**: Modify the value of RIR network attributes by clicking the **Value** field of an attribute and entering a new value. You can add a new RIR network attribute by clicking the Add icon and selecting an attribute from the drop-down list. You can also select any optional attributes and click the Delete icon to delete them. For information about RIR network attributes, see [RIR Network Attributes](#).

You can enter up to 256 characters for all RIR network attributes, unless otherwise noted.

**Preview RIR Submissions**: Click this to view the updates before the appliance submits them to the RIPE database. This button is enabled only when the registration action is **Create**, **Modify**, or **Delete**, and the **Do not**

**update registrations** checkbox is not selected. For more information, see [Previewing Registration Updates](#).

To schedule this task, click the Schedule icon at the top of the wizard. In the Schedule Change panel, click **Later**, and then specify a date, time, and time zone.

3. Save the configuration.

## Deleting RIR Networks

When you delete an RIR network or network container, the appliance moves it to the Recycle Bin, if enabled. You must enter the reason for deleting the RIR network or network container and indicate whether you want to send the deletion update to RIPE. You can delete multiple networks at the same time.

To delete an RIR network or network container:

1. From the **Data Management** tab, select the **IPAM** tab -> *network* checkbox, and then click the Delete icon. You can choose to delete the network immediately or schedule its deletion.
2. In the *Delete Confirmation (IPv4 or IPv6 Network)* or *Schedule Deletion* dialog box, complete the following:
  - **Justification:** Enter the reason for deleting this network.
  - **Do not update registrations:** Select this checkbox if you do not want the appliance to submit updates to RIPE.
3. Optionally, you can click **Preview RIR Submissions** to view the RIR network information before deleting the network. Grid Manager displays the preview data in a separate browser window. For information, see [Previewing Registration Updates](#)
4. Click **Yes**. If you are scheduling a deletion, enter the data and time for execution, and then click **Schedule Deletion**.

## Managing RIR Attributes

Before you can successfully submit RIR data updates, you must ensure that all RIR required attributes contain valid values that can be mapped to data in the RIPE database. The appliance does not validate data with the RIPE database before you submit your updates. The appliance also does not synchronize data from the database.



### Note

RIPE does not support UTF-8 data in the **Description** and **Remarks** fields. After an upgrade, the NIOS appliance keeps the UTF-8 data in these fields. However, if you want to modify these fields after the upgrade, you must remove the UTF-8 data before you can save the changes.

When you enter a value for the following RIR attributes that cannot be mapped to a valid reference in the RIPE database, updates to the RIR database will fail. However, these values will still be displayed in the IPv4 or IPv6 network or network container panels of Grid Manager.

- RIPE Routes Maintainer
- RIPE Lower Level Maintainer
- RIPE Reverse Domain Maintainer
- RIPE Admin Contact
- RIPE Technical Contact
- RIPE Computer Security Incident Response Team

You can add multiple values for certain RIR attributes. When you add multiple values of the same attribute, the appliance groups the values in the order they are listed in the attribute table. You can also reorder the RIR attributes using the up and down arrows in the attribute tables.

## RIR Organizational Attributes

The following table lists RIR organizational attributes, the format you must use to enter values, and whether they are required or optional.

Organizational Attribute	Corresponding RIPE Attribute	Description and Format	Required/ Optional
<b>RIPE Description</b>	descr	Enter a short description about the organization.	Optional
<b>RIPE Country</b>	country	From the drop-down menu, select the country name, followed by the two-letter ISO 3166 country code, of the country or area within the RIPE NCC service region or through Local Internet Registries.	Required
<b>RIPE Admin Contact</b>	admin-c	Enter the name of the on-site admin contact for the organization. Enter the name in this format: Start with two to four optional characters, followed by up to six optional digits, and then follow by a source specification. The first digit cannot be "0". The source specification starts with "-" followed by the source name that contains up to nine characters in length.	Required
<b>RIPE Technical Contact</b>	tech-c	Enter the name of the technical contact for the organization. Enter the name in this format: Start with two to four optional characters, followed by up to six optional digits, and then follow by a source specification. The first digit cannot be "0". The source specification starts with "-" followed by the source name that contains up to nine characters in length.	Required
<b>RIPE Remarks</b>	remarks	Enter remarks about the organization.	Optional
<b>RIPE Notify</b>	notify	Enter the email address to which notifications of changes to the organization will be sent.	Optional
<b>RIPE Registry Source</b>	source	From the drop-down list, select the registry at which the organization is registered. The default is RIPE. Select <b>RIPE</b> for the RIPE database, which is the authoritative database. Select <b>TEST</b> for the RIPE TEST database that operates in the same way as the RIPE database but contains only test data. Note that test data is cleaned out at the start of each month and a predetermined set of basic objects is re-inserted. You can use the RIPE TEST database to learn how to update the database and try out special scenarios. The RIPE TEST database has fewer restrictions which allows you to create encompassing or parent objects you may need for testing.	Optional
<b>RIPE Organization Type</b>	org-type	<p>From the drop-down list, select one of the following organization type:</p> <ul style="list-style-type: none"> <li>• <b>IANA</b> for Internet Assigned Numbers Authority</li> <li>• <b>RIR</b> for Regional Internet Registry</li> <li>• <b>NIR</b> for National Internet Registry</li> <li>• <b>LIR</b> for Local Internet Registry</li> <li>• <b>WHITEPAGES</b> for special industry people</li> <li>• <b>DIRECT_ASSIGNMENT</b> for direct contract with RIPE NCC</li> <li>• <b>OTHER</b> for all other organizations</li> </ul> <hr/> <p><b>Note:</b> Only the RIPE database admin can set the organization type, and there are no NIRs in the RIPE NCC service region.</p> <hr/>	Optional
<b>RIPE Address</b>	address	Enter the organization address.	Optional

Organizational Attribute	Corresponding RIPE Attribute	Description and Format	Required/ Optional
<b>RIPE Phone Number</b>	phone	Enter the organization phone number in numeric format starting with the + character, followed by the country code, area code, and the phone number. For example, you can enter <b>+18089991000</b> . You can also use one of the following formats: <ul style="list-style-type: none"> <li>• '+' &lt;integer-list&gt;</li> <li>• '+' &lt;integer-list&gt; "(" &lt;integer-list&gt; ")" &lt;integer-list&gt;</li> <li>• '+' &lt;integer-list&gt; ext. &lt;integer list&gt;</li> <li>• '+' &lt;integer-list&gt; "(" integer list ")" &lt;integer-list&gt; ext. &lt;integer-list&gt;</li> </ul>	Optional
<b>RIPE Fax Number</b>	fax-no	Enter the organization fax number in numeric format starting with the + character, followed by the country code, area code, and the fax number. For example, you can enter <b>+16052529000</b> . You can also use one of the following formats: <ul style="list-style-type: none"> <li>• '+' &lt;integer-list&gt;</li> <li>• '+' &lt;integer-list&gt; "(" &lt;integer-list&gt; ")" &lt;integer-list&gt;</li> <li>• '+' &lt;integer-list&gt; ext. &lt;integer list&gt;</li> <li>• '+' &lt;integer-list&gt; "(" integer list ")" &lt;integer-list&gt; ext. &lt;integer-list&gt;</li> </ul>	Optional
<b>RIPE Email</b>	email	Enter the organization email address.	Required
<b>RIPE Abuse Mailbox</b>	abuse-mailbox	Enter the email address to which abuse complaints are sent.	Optional
<b>RIPE Reference Notify</b>	ref-nfy	Enter the email address to which notifications are sent when a reference to the organization object is added or removed.	Optional

## RIR Network Attributes

When you create or edit an RIR associated network, ensure that you enter valid values for the RIR network attributes. The following table lists RIR network attributes, the format you must use to enter values, and whether they are required or optional:

Network Attributes	Corresponding RIPE Attribute	Descriptions and Formats	Required/ Optional
<b>RIPE Admin Contact</b>	admin-c	The name of the on-site admin contact for the network address. This attribute is populated from the organizational attribute. You can modify the value in this format: Start with two to four optional characters, followed by up to six optional digits, and then follow by a source specification. The first digit cannot be "0". The source specification starts with "-" followed by the source name that contains up to nine characters in length.	Required
<b>RIPE Computer Security Incident Response Team</b>	mnt-irt	The name of the Computer Security Incident Response Team (CSIRT) that handles security incidents for the network address. You can enter the value in this format: Use letters, digits, the character underscore "_", and the character hyphen "-". The value must start with "irt-", and the last character of a name must be a letter or a digit. You must enter a minimum of five characters.	Optional
<b>RIPE Country</b>	country	The two-letter ISO 3166 country code of the country within the RIPE NCC service region or through Local Internet Registries. This attribute is populated from the organizational attribute. You can select a different country code from the drop-down list.	Required

Network Attributes	Corresponding RIPE Attribute	Descriptions and Formats	Required/Optional
<b>RIPE Description</b>	descr	Enter a short description about the network.	Required
<b>RIPE IPv4 Status</b>	status	The status of the IPv4 network address. From the drop-down list, select one of the following status: ALLOCATED PA ALLOCATED PI ALLOCATED UNSPECIFIED LIR-PARTITIONED PA LIR-PARTITIONED PI SUB-ALLOCATED PA ASSIGNED PA ASSIGNED PI ASSIGNED ANYCAST EARLY-REGISTRATION NOT-SET	Required
<b>RIPE IPv6 Status</b>	status	The status of the IPv6 network address. From the drop-down list, select one of the following: ALLOCATED-BY-RIR ALLOCATED-BY-LIR ASSIGNED ASSIGNED ANYCAST ASSIGNED PI	Required
<b>RIPE Lower Level Maintainer</b>	mnt-lower	Enter the name of the registered maintainer for hierarchical authorization purposes. This can protect the creation of networks directly (one level) below in the hierarchy of a network container or another network. The authentication method of the maintainer will be used upon creation of any network directly below the network that contains the "mnt-lower:" attribute. Enter the maintainer name in this format: Use letters, digits, the character underscore "_", and the character hyphen "-". The first character must be a letter, and the last character must be a letter or a digit. You cannot use the following words (they are reserved by RPSL): any, as-any, rs-any, peer, as, and, or, not, atomic, from, to, at, action, accept, announce, except, refine, networks, into, inbound, outbound. Also note the following: Names starting with certain prefixes are reserved for certain object types. Names starting with "as-" are reserved for as set names. Names starting with "rs-" are reserved for route set names. Names starting with "rtrs-" are reserved for router set names. Names starting with "ftr-" are reserved for filter set names. Names starting with "prng-" are reserved for peering set names. Names starting with "irt-" are reserved for irt names.	Optional
<b>RIPE Network Name</b>	netname	The name of the IP address range. You can enter up to 80 characters. Enter the network name in this format: Use letters, digits, the character underscore "_", and the character hyphen "-". The first character must be a letter, and the last character must be a letter or a digit.	Required
<b>RIPE Notify</b>	notify	Enter the email address to which notifications of changes to the object must be sent.	Optional
<b>RIPE Registry Source</b>	source	From the drop-down list, select the registry at which the organization is registered. The default is <b>RIPE</b> . Select <b>RIPE</b> for the <b>RIPE</b> database, which is the authoritative database. Select <b>TEST</b> for the <b>RIPE TEST</b> database that operates in the same way as the <b>RIPE</b> database but contains only test data. Note that test data is cleaned out at the start of each month and a predetermined set of basic objects is re-inserted. You can use the <b>RIPE TEST</b> database to learn how to update the database and try out special scenarios. The <b>RIPE TEST</b> database has fewer restrictions which allows you to create encompassing or parent objects you may need for testing.	Required
<b>RIPE Remarks</b>	remarks	Enter remarks about the network.	Optional

Network Attributes	Corresponding RIPE Attribute	Descriptions and Formats	Required/Optional
<b>RIPE Reverse Domain Maintainer</b>	mnt-domains	Enter the name of a registered maintainer used for reverse domain authorization. This can protect domain objects. The authentication method of this maintainer will be used for any encompassing reverse domain object. Enter the maintainer name in this format: You can use letters, digits, the character underscore "_", and the character hyphen "-". The first character must be a letter, and the last character must be a letter or a digit. You cannot use the following words (they are reserved by RPSL): any, as-any, rs-any, peer, as, and, or, not, atomic, from, to, at, action, accept, announce, except, refine, networks, into, inbound, outbound. Also note the following: Names starting with certain prefixes are reserved for certain object types. Names starting with "as-" are reserved for as set names. Names starting with "rs-" are reserved for route set names. Names starting with "rtrs-" are reserved for router set names. Names starting with "fltr-" are reserved for filter set names. Names starting with "prng-" are reserved for peering set names. Names starting with "irt-" are reserved for irt names.	Optional
<b>RIPE Routes Maintainer</b>	mnt-routes	This attribute references a maintainer that is used in determining authorization for the creation of route objects. Enter the name in this format: Start with the reference to the maintainer, followed by an optional list of prefix ranges inside of curly brackets or the keyword "ANY". The default, when no additional set items are specified, is "ANY". For more information, refer to RFC-2622. Example: <mnt-name> [ { list of <address-prefix-range> }   ANY ].	Optional
<b>RIPE Technical Contact</b>	tech-c	The name of the technical contact for the network. Enter the name in this format: Start with two to four optional characters, followed by up to six optional digits, and then follow by a source specification. The first digit cannot be "0". The source specification starts with "-" followed by the source name that contains up to nine characters in length.	Required

## Monitoring RIR Data

You can view RIR organizations and networks you added to NIOS through Grid Manager. The appliances sends SNMP traps and email notifications about registration updates. It also logs RIR events in the Infoblox syslog. Note that sometimes due to network timeout from RIPE, your registration updates may fail.

You can do the following to monitor RIR data:

- View RIR update events in the syslog, as described in [Viewing the Syslog](#).
- View RIR organizations, as described in [Viewing RIR Organizations](#) below.
- View RIR IPv4 and IPv6 network containers and networks, as described in [IPAM Home](#) and [Viewing Networks](#).
- Preview RIR updates before submitting them to RIPE, as described in [Previewing Registration Updates](#) below.

## Viewing RIR Organizations

You can view the list of RIR organizations that have received address allocation and the ones you have added associated networks.

To view RIR organizations:

1. From the **Administration** tab, select the **RIR** tab.
2. Grid Manager displays the following information for each RIR organization:
  - **Organization ID:** The RIR organization ID.
  - **RIR:** The RIR that allocates the address block to the organization.
  - **Maintainer:** The name of the maintainer for the organization.

You can also select **Organization Name** and RIR organizational attributes for display.

You can do the following in this tab:

- Modify some of the data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click Save to save the changes. Note that some fields are read only.
- Sort the data in ascending or descending order by column.
- Select an organization and click the Edit icon to modify data, or click the Delete icon to delete it.



- Click the Permissions icon to configure permissions for the admin account.
- Use filters and the Go to function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the Go to field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information about quick filters, see [About the Grid Manager Interface](#).
- Print and export the data in this tab.

## Previewing Registration Updates

Before the appliance submits RIR updates to RIPE, you can preview the data in Grid Manager. The appliance uses the email template when displaying preview data in a separate browser window.

Preview data includes the subject line for the email, followed by the inetnum or inet6num template for the network and other associated data, such as network name, organization name, and others. When there are multiple operations involved, such as deleting multiple networks, the preview data includes a separate subject line for each operation. You can preview registration updates when you create a new RIR network. In the Add Networks wizard, click **Preview RIR Submissions** in the wizard. For information about how to create or assign RIR networks, see [Adding and Assigning RIR Networks](#).

Following is a sample preview for a network creation request:

```
Subject: CREATE inetnum 100.200.0.0 - 100.200.255.255 KEYWORDS: NEW inetnum:
100.200.0.0 - 100.200.255.255
netname: corpxyz_network
descr: RIR network for corpxyz. status: ASSIGNED PA
org: ORG-MC1-TEST
country: US source: TEST
changed: jdoe@corpxyz.com 20120809 notify: jdoe@corpxyz.com
admin-c: NP1-TEST tech-c: NP1-TEST mnt-by: JohnDoe password: *** password:
***
```

Following is a sample preview for a network modification request:

```
Subect: MODIFY inetnum 100.200.0.0 - 100.200.255.255 KEYWORDS: inetnum:
100.200.0.0 - 100.200.255.255
netname: corpxyz_network
descr: RIR network for corpxyz. status: ASSIGNED PA
org: ORG-MC1-TEST
country: US source: TEST
changed: jdoe@corpxyz.com 20120809 notify: jdoe@corpxyz.com
admin-c: NP1-TEST tech-c: NP1-TEST mnt-by: JohnDoe password: ***
```

Following is a sample preview for deleting multiple networks:

```
Subect: DELETE inetnum 100.200.0.0 - 100.200.255.255 KEYWORDS: inetnum:
100.200.0.0 - 100.200.255.255
netname: corpxyz_network
descr: RIR network for corpxyz. status: ASSIGNED PA
org: ORG-MC1-TEST
```

```
country: US source: TEST
```

```
changed: jdoe@corpxyz.com 20120809 notify: jdoe@corpxyz.com
```

```
admin-c: NP1-TEST tech-c: NP1-TEST mnt-by: JohnDoe password: ***
```

```
delete: Removed network.
```

```
Subject: DELETE inetnum 100.300.0.0 - 100.300.255.255 KEYWORDS: inetnum:  
100.300.0.0 - 100.300.255.255
```

```
netname: corp200_network
```

```
descr: RIR network for Corp200. status: ASSIGNED PA
```

```
org: ORG-MC1-TEST
```

```
country: US source: TEST
```

```
changed: jsmith@corp200.com 20120809 notify: jsmith@corp200.com
```

```
admin-c: NP1-TEST tech-c: NP1-TEST mnt-by: JohnSmith password: ***
```

```
delete: Removed network.
```

## Requirements and Permissions

To manage RIR allocated addresses, organizations, and network utilization that contain RIR assignments, you must first enable support for RIR registration updates, and then configure the RIR communication method. Note that once you have enabled support for RIR registration updates, settings and fields that are relevant to this feature are enabled in Grid Manager. You do not need a special license to use this feature.

Only superusers can create, modify, and delete RIR organizations. Limited-access users can manage RIR allocated address blocks if they have the required permissions to the objects.

To view and manage RIR related data, admins must have permissions to the applicable resources. For example, to view RIR networks, admins must have read-only permission to the networks; and to edit them, admins must have read/write permission to them. For more information about admin permissions, see [About Administrative Permissions](#).

## IP Address Management

IPAM (IP Address Management) is the allocation, administration, reporting, and tracking of IP addresses, network devices, and their associated data. This section provides information about IPAM and how to use the Infoblox tools to perform IPAM tasks and manage your entire IP network. It includes the following topics:

- [Managing IP Addresses](#)
- [IP Discovery and vDiscovery](#)
- [Infoblox Network Insight](#)

## Managing IP Addresses

This section describes how to manage your networks and IP addresses through the Infoblox IPAM (IP Address Management) implementation. It contains the following topics:

- [About IP Address Management](#)
- [About Host Records](#)
- [Managing IPv4 Networks](#)
- [Viewing and Managing IPv4 Addresses](#)
- [Managing IPv6 Networks](#)

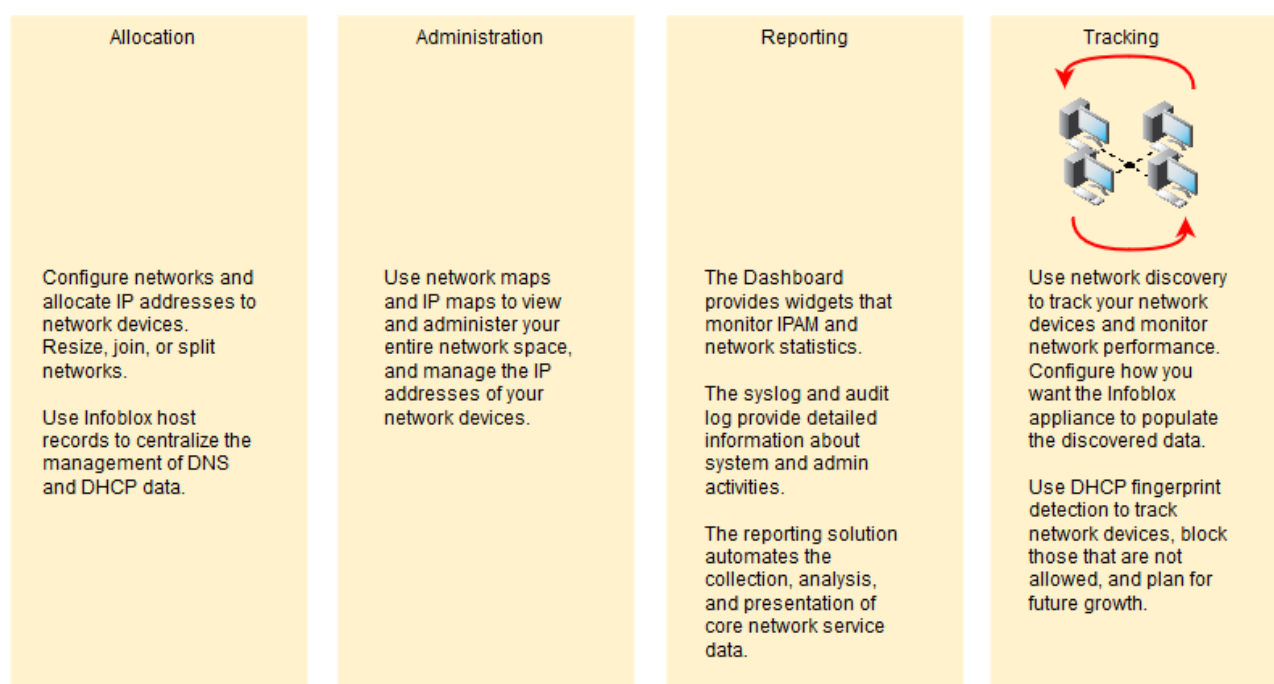
- [Viewing IPv6 Data](#)
- [Managing IPv4 and IPv6 Addresses](#)
- [Configuring Thresholds for IPAM Utilization](#)
- [Viewing Identity Mapping Information](#)

## About IP Address Management

IPAM is the allocation, administration, reporting, and tracking of public and private IP spaces, network devices, and their associated data. It comprises the deployment of DNS and DHCP services and the monitoring of network devices and performance to ensure data integrity and security of your networks.

The Infoblox IPAM implementation is a feature-rich and easy-to-use solution that encompasses support for IPv4, IPv6, network discovery, and automated monitoring. Infoblox IPAM provides tools that integrate the allocation, administration, reporting, and tracking of your entire network space. The below figure highlights the Infoblox IPAM implementation.

### Infoblox IPAM Features



You can perform the following IPAM tasks to effectively manage and control your network:

- Create host records. Host records integrate the DNS records and DHCP data of a network device. You can use a host record to manage a network device from one central point. For more information, see [About Host Records](#).
- Create IPv4 and IPv6 networks. For information, see [About Network Containers](#).
- Discover devices in IPv4 and IPv6 networks and other Objects created in IPAM. For more information about Adding IPv4 and IPv6 Network Containers and Networks and Adding Host Records, see [About Host Records](#). View your IPv4 network and address utilization in a graphical mode. For information, see [IPv4 Network Map](#) and [IP Map](#).
- When necessary, resize, join, or split networks. For information about Resizing IPv4 Networks, Splitting IPv4 Networks into Subnets, Joining IPv4 Networks see, [Managing IPv4 Networks](#). For information about Splitting IPv6 Networks into Subnets, and Joining IPv6 Networks, see [Managing IPv6 Networks](#).
- Manage IPv4 and IPv6 address data. For information, see [Viewing and Managing IPv4 Addresses](#), [Viewing IPv6 Data](#), and [Managing IPv4 and IPv6 Addresses](#).

- Add and manage DNS resource records associated with IP addresses. For information, see [Configuring DNS Resource Records](#).
- Monitor your core service network data using the Dashboard, audit log, syslog, and reports. For information, see [Infoblox Reporting and Analytics](#).
- Discover and track network devices. For information, see [IP Discovery and vDiscovery](#) and [DHCP Fingerprint Detection](#).

## About Host Records

Host records provide a unique approach to the management of DNS, DHCP, and IPAM data. By using host records, you can manage multiple DNS records and DHCP and IPAM data collectively, as one object on the appliance.

When you create a host record, you are specifying the name-to-address and address-to-name mappings for the IP address that you assign to the host. The Infoblox DNS server then uses this data to respond to DNS queries for the host. When the server receives a name-to-address query, it responds with an A record for an IPv4 host or an AAAA record for an IPv6 host that contains the data from the host record. Likewise, when it receives an address-to-name query for the host, the appliance responds with a PTR record that contains data from the host record.



### Note

The appliance cannot respond if there is no PTR record and a PTR record is not created if there is no corresponding reverse-mapping zone.

Additionally, if you specify an alias in the host record, the appliance uses this data as a CNAME record to respond to queries with the alias. It maps the alias to the canonical name and sends back a response with the canonical name and IP address of the host. Thus, a single host record is equivalent to creating A, PTR, and CNAME resource records for an IPv4 address and AAAA and PTR records for an IPv6 address. The appliance supports IDNs for a host record. You can specify alias and domain names in the native character set. For information about IDN support, see .

Hosts also support prefix delegation for IPv6. For example, you can specify an IPv6 prefix in the host record of a router. The router then advertises this prefix on one of its interfaces, so hosts that connect to the interface can generate their IP addresses, using the stateless autoconfiguration mechanism defined in *RFC 2462, IPv6 Stateless Autoconfiguration*. In addition, if the Infoblox DHCP server manages the IP address assigned to the host, the server uses it as a fixed address record as well. The DHCP server assigns the IP address to the host when it receives a DHCP request with the matching MAC address or DUID. Its response includes configuration information, and any DHCP options defined for the host or inherited from the network to which the fixed address belongs. You can also assign multiple IPv4 and IPv6 addresses to a host, as described in [Assigning Multiple IP Addresses to a Host](#) below.

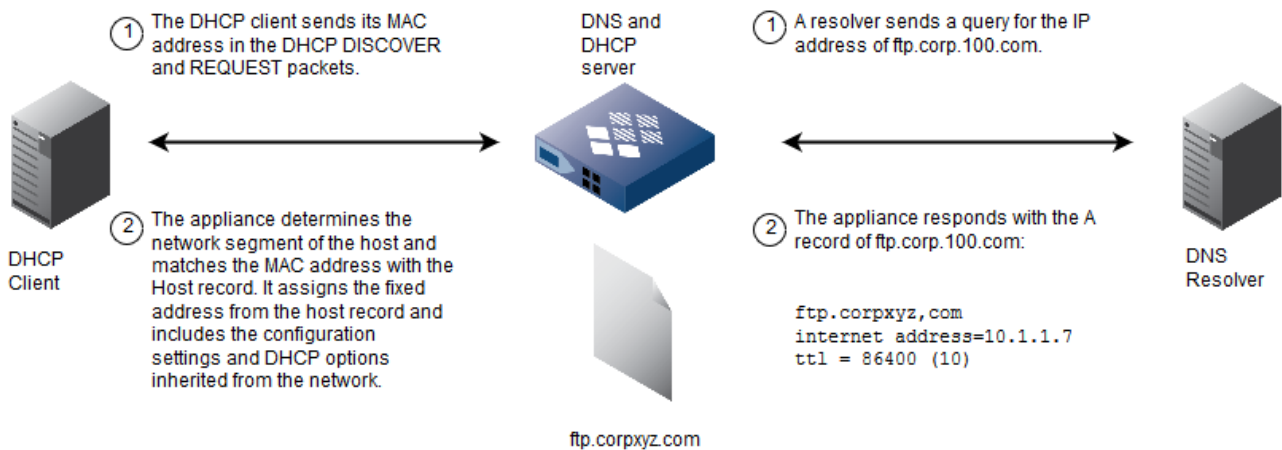
You can copy an existing host record and turn it into a new one. When you copy a host record, other than the new host name and IP address, all DHCP and IPAM configuration including the MAC address and extensible attributes apply to the new record. You can also modify information, except for the host name and IP address, of an existing host record. For information about how to copy or modify a host record, see [Copying and Modifying Host Records](#) below. Note that you can also modify an IPv4 host record and turn it into a IPv4 reservation. For information, see [Configuring IPv4 Reservations](#).

You can execute immediate discovery on a host record. This simple setting enables you to determine the precise type of device that is associated with the host, along with its IP addresses, its name and other information.

You can define extensible attributes for a host record to further describe the device. You can include information such as its location and owner for IP address management purposes. For information about extensible attributes, see [About Extensible Attributes](#).

The below figure illustrates how the appliance uses the host record for both DHCP and DNS.

*Using the Host Record for DHCP and DNS*

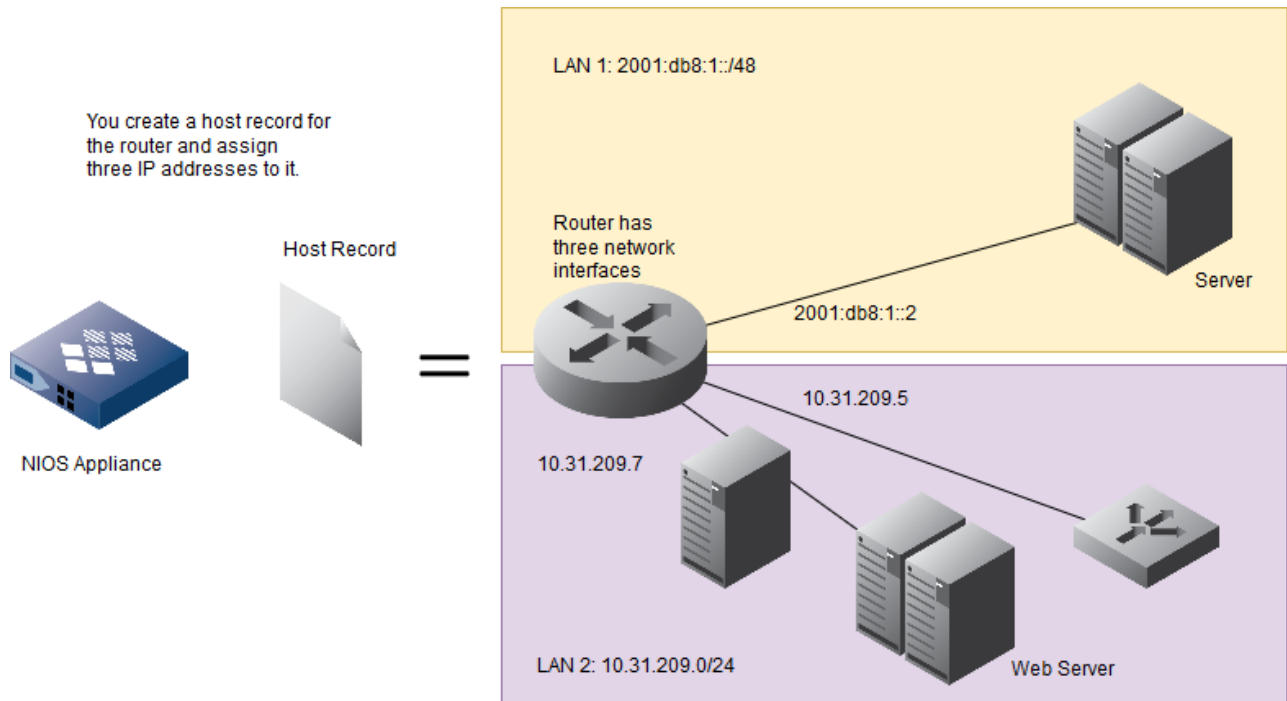


Note that If the zone of the host record is associated with networks, the IP addresses must belong to the associated networks. For example, if the host record is in the [corpxyz.com](#) zone, which is associated with 10.1.0.0/16 network, then the IP addresses of the host record must belong to the 10.1.0.0/16 network. For information about associating zones and networks, see [Associating Networks with Zones](#).

### Assigning Multiple IP Addresses to a Host

You can assign multiple IPv4 and IPv6 addresses to a host depending on the function of the device. For example, you can create a host record for a router that supports three network interfaces in two different networks, and assign IP addresses to each interface, as illustrated in the figure [Assigning Multiple IP Addresses to one Host Record](#) below. When the DNS server responds to DNS queries for the host, it includes an A or AAAA record for each IP address. In addition, if the IP addresses belong to different networks, they can have different DHCP configurations and options. As shown in the figure [Assigning Multiple IP Addresses to one Host Record](#) below, the configuration information and DHCP options of the interface with the IPv6 address 2001:db8:1::2 may be different from the other two interfaces, 10.31.209.5 and 10.31.209.7, because it is in a different network.

*Assigning Multiple IP Addresses to one Host Record*



## Adding Host Records

You can add host records from the Toolbar of the **IPAM**, **DHCP** and **DNS** tabs of the **Data Management** tab and from the Tasks Dashboard. For information about the Tasks Dashboard, see [Tasks Dashboard](#).

When you create a host record, you must specify its zone and at least one IP address. If the zone of the host record is associated with one or more networks, the IP addresses must belong to one of the associated networks. If a zone of a host record contains IDNs, the appliance displays the zone name in the native character set.

For Cloud Network Automation, you can create host records within a delegated network view only when you enable DNS for the host record.

To add a host from the **Data Management** tab:

1. From the **IPAM**, **DHCP** or **DNS** tab of the **Data Management** tab, expand the Toolbar.
2. Click **Add** and select the option to add a host from the drop-down menu.
3. In the first page of the *Add Host* wizard, do the following:
  - **Name:** If Grid Manager displays a zone name, enter the host name here. The displayed zone name can either be the last selected zone or the zone from which you are adding the host. If no zone name appears or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter a unique name for the host. The name you enter is prefixed to the DNS zone name that is displayed, and the complete name becomes the FQDN (fully qualified domain name) of the host. For example, if the zone name displayed is corpxyz.com and you enter admin, then the FQDN is admin.corpxyz.com.
  - **Enable in DNS:** This is selected by default. It enables DNS service for the host. If you clear this checkbox, DNS does not serve this host and you cannot assign it to a zone.
  - **Protected:** This is selected by default. To mark the record as protected in order to restrict DDNS updates to this record, select the Protected checkbox. This applies to both static and dynamic records. Note that if you select to protect the record, ensure that you also select the Prevent dynamic updates to RRsets containing protected records checkbox in the advanced updates properties of the Grid, view, zone, or Standalone appliance.

Alternatively, you can protect records by selecting them, individually or in bulk, in the Resource Records Viewer and clicking Protect Records -> Enable Protection in the Toolbar.

- **DNS View:** Displays the DNS view for the host record. This appears only when you enable the host record in DNS.
- **Host Name Policy:** Displays the host name policy of the selected zone. This appears only when you enable the host record in DNS.
- **RRset Order:** Select one of the following RRset orders that the appliance uses to return A and AAAA records of the host. This checkbox appears only when you have enabled the configuration of RRset order for the Grid and there are multiple IP addresses in this host record. For information about how to enable this feature, see [Enabling the Configuration of RRset Orders](#).
  - **Cyclic:** The records are returned in a round robin pattern. This is the default.
  - **Fixed:** The records are returned in the order you specify in this host record. When you select this checkbox, the appliance displays up and down arrows next to the IPv4 and IPv6 address tables. You can use these arrows to reorder the address list. The appliance returns the A and AAAA records of this host based on the order you define in the address tables.
  - **Random:** The records are returned in a random order. Note that when you specify **Fixed** as the RRset order, the appliance places the resource records as follows:
    - A and AAAA records of the host in the fixed order you specify in the address tables. Note that the order of the returned A and AAAA records are independent of each other.
    - Other A and AAAA records in an undefined order.
    - Other record types in the default cyclic order.For more information about RRset order, see [Enabling the Configuration of RRset Orders](#).
- In the **IPv4 Addresses** and **IPv6 Addresses** sections, specify the IP addresses of the host record. Click the Add icon do one of the following:
  - Select **Next Available IP Address** to retrieve the next available IP address in a network. Infoblox recommends this option to ensure that you assign an IP address from the appropriate network. If the host record is in a zone that has one associated network, Grid Manager retrieves the next available IP address in that network. If the host record is in zone that has multiple associated networks, the *Network Selector* dialog



box lists the associated networks. If the zone has no network associations, the *Network Selector* dialog box lists the available networks. When you select a network, Grid Manager retrieves the next available IP address in that network.

If you want to enter a link-local IPv6 address, you must enter an IPv4 address and the host MAC address first, and then click the Add (+) icon again to enter the link-local IPv6 address. When you select the link-local IPv6 address, the MAC address is automatically filled in. For information, see [Understanding DNS for IPv6](#).

Optionally, you can delete an IP address from the host by selecting an IP address in the table and clicking the Delete icon.

or

- Select **Add Address** to enter an IPv4 or IPv6 address. You can also enter an IPv6 prefix. Note that when you use this option, you could specify an IP address from a network that has not yet been defined. To avoid this, use the **Next Available IP Address** option instead.
- **MAC Address**: For an IPv4 address, enter the MAC address of the network device associated with this host IP address. Note that you must enter a MAC address if DHCP is enabled for the host IP address.

or

- **DUID**: For an IPv6 address, enter the DHCP Unique Identifier (DUID) of the network device associated with this host IP address. Note that you must enter a DUID if DHCP is enabled for an IPv6 host address.
- **DHCP**: Select this to enable the DHCP services to manage the host IP address. If you do not select this option, the host IP address is not managed by the DHCP server.
- **Comment**: Optionally, enter additional information about the host record.
- **Disable**: Select this option to temporarily disable the host record. For example, you might want to disable a host when you need to update the network device.

The **Cloud** section appears when the Cloud Network Automation license is installed on the Grid Master. For information, see [Deploying Cloud Network Automation](#). This section displays the following information:

- **Cloud Usage**: This field indicates whether this object is associated with any specific cloud extensible attributes or within a scope of delegation. It can be one of the following:
  - **Cloud from adapter**: Indicates that this object has been created by a cloud adapter and it may or may not be within a scope of delegation at the moment.
  - **Cloud from delegation**: Indicates that this object is within the scope of delegation or the object itself defines a scope of authority delegation, and it is not created by a cloud adapter.
  - **Used by cloud**: Indicates that this network or network container is associated with the extensible attribute **Is External** or **Is Shared** and the value is set to True, which implies the network is a private or shared network managed by the CMP, and it is not **Cloud from adapter** or **Cloud from delegation**.
  - **Non-cloud**: The object is a regular NIOS object and is not within the scope of any authority delegation nor is it associated with any of these extensible attributes: **Cloud API Owned**, **Is External** or **Is Shared**. NIOS admin users can modify this object based on their permissions.
- **Owned By**: A cloud object can be owned by the Grid Master or the cloud adapter. When the object is created by the Grid Master, this shows **Grid**. If the object is created by the cloud adapter, this shows **Adapter**.  
Delegate authority from the Grid Master
- **Delegate To**: This field indicates whether the authority for the object you want to create has already been delegated. If so, it displays the name of the delegation.

4. (*Applies only to Network Insight*) In the current Wizard step, you can optionally define the following identification values and settings for the new object's port reservation:

- Choose the Device Type: Router, Switch-Router, Switch, MSFT (Microsoft) Server, **NetMRI**, **NIOS**, **VNIOS**, or ESX (VMware) Server.

The values on this page are not required for defining the actual port reservation in a later wizard step.

- Choose the Device Vendor: Cisco, Juniper, Aruba, Dell, **Infoblox**, or HP.
- You can also enter a Location and a Description. These values are advisory and not required for configuration.

After you define this group of settings, you will still need to define a device port reservation.

5. (Applies only with Network Insight) Click **Next** to initiate or disable discovery of the new host.
  - Choose either **Exclude from Network Discovery** or **Enable Immediate Discovery**. If you choose to Exclude, discovery will not execute on the host. If you choose **Enable Immediate Discovery**, discovery will execute on the host after you save your settings. You may also choose to leave both options disabled.
  - By default, the new host inherits its SNMP credentials from those defined at the grid level. Should you wish to override them for a local set of credentials, check the **Override Credentials** checkbox and select the **SNMPv1/SNMPv2** or **SNMPv3** option and enter the locally used credentials. For more information, see the sections [Configuring SNMP1/v2 Credentials for Polling](#) and [Configuring SNMPv3 Properties](#) for a complete description of SNMP credentials for discovery. (You can also test SNMP credentials to ensure they work before use.)
  - For the new object, you can check the **Override CLI Credentials** checkbox to override the inherited set of CLI credentials taken from the Grid level. This set of credentials may be used for the device that is directly associated with the new object (in this case, a Host) in its port reservation.
  - You can also click **Test CLI Credentials** to enter and test a set of CLI login credentials against a device based on its IP address.  
Port control operations require CLI credentials for the involved devices. (If you are not using port control for the new object, usage of CLI credentials is optional.) Because some IPAM and DHCP objects will use port control features as part of object creation, CLI credentials are automatically leveraged as part of discovery. Ensure you have the correct sets of CLI credentials for devices in your network. For more information, the section [Configuring CLI Discovery Properties](#).
  - SSH is the default for CLI operations. Check the **Allow Telnet** checkbox if you know the device involved in the object assignment may support Telnet but may not support SSH, or if you want Telnet as an option.
6. (Applies only with Network Insight) Click **Next** to define switch port connectivity for the device that will be associated with the new host record. *This step is optional and not required for creating the new host record.* This feature set is also termed port control in Grid Manager. The device to which the new host record will be associated should already be discovered and managed from Grid Manager.
  - Begin by checking the Reserve Port checkbox. Note that reserving a switch port does not guarantee its availability.  
Optionally, you can skip connecting port configuration by clicking Next.  
Click the **Clear** button to remove the selected device from the configuration.
  - Click the Select Device button to choose the device for which the port reservation will be associated. You should know the identity of the device to whose interface the new object will be associated before taking this step. For more information, see the section [Using the Device Selector](#).
  - After choosing the device, choose the Interface with which the port reservation will be bound. The drop-down list shows only interfaces that are most recently found to be available by Grid Manager during the last discovery cycle. This list will not include any ports that are Administratively Up and Operationally Up or that are otherwise already assigned to other networks or objects.
  - The Wizard page also shows a list of any VLANs that are currently configured in the chosen device (**The following VLANs are configured**). This Wizard page allows only the assignment of an existing VLAN in the chosen device to the new port reservation.
  - Check the Configure Port checkbox to define specific port control settings for the port reservation.
  - Choose the Data VLAN and/or the Voice VLAN settings you may need for the port assignment.  
Depending on the selected device, you may or may not be able to apply VLAN settings.
  - Set the Admin Status to Up if you need to activate the port after assignment in the current task.  
All port control operations require CLI credentials to be entered into Grid Manager. Because some IPAM and DHCP objects will use port control features as part of object creation, CLI credentials are automatically leveraged as part of discovery and configuration of port configurations such as Admin Up/Down status.  
Ensure you have the correct sets of CLI credentials for devices in your network.
  - Enter a Description for the port assignment. Infoblox recommends doing so to help other technicians to recognize the port assignment task.
7. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).
8. As the final step in the Add Host wizard, you define when Grid Manager creates the new object by scheduling it. As a separate task, you also schedule when the associated Port Configuration task executes.
  - To create the new Host and its associated port reservation immediately, select Now. The port control event is automatically synchronized to take place at the same time as the activation of the new host.
  - You can choose to have Grid Manager execute the port reservation at the same time as the host object creation. To do so, select **At same time as Host**.



- You can have Grid Manager execute the port reservation at a later time by selecting Later. Choose a **Selected time** by entering or selecting a Start Date (click the calendar icon to choose a calendar date) and a Start Time, and choose a Time Zone.
9. Choose one of the following from the **Save & ...** drop-down button menu:
    - Click **Save & Close** to add the Host object and close the wizard (this is the default).
    - Click **Save & Edit** to add the Host object and launch the editor.
    - Click **Save & New** to add the Host object and launch the wizard again to add another Host object.

## Copying and Modifying Host Records

You can create a new host record by copying an existing one. When you copy a host record, other than the new host name and IP address, all DHCP and IPAM configuration including the MAC address and extensible attributes apply to the new record. You can also modify information, except for the host name and IP address, of an existing host record.

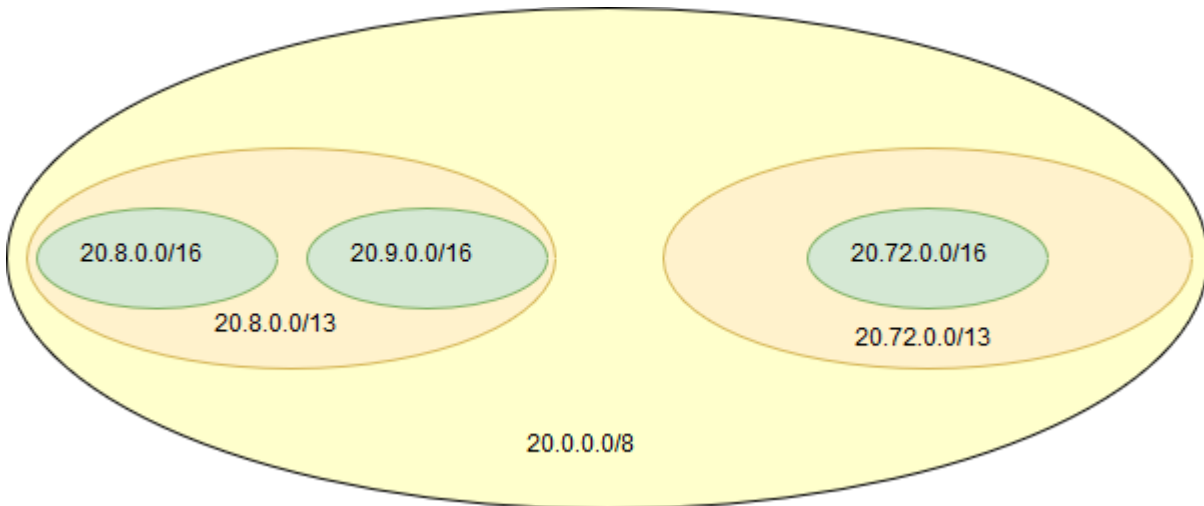
To copy or modify an existing host record:

1. From the **Data Management** tab, select the **IPAM, DHCP, or DNS** tab.
2. In the selected tab, search for or navigate to the host record that you want to copy or modify.
3. Do one of the following:
  - To copy a host record, select the record and expand the Toolbar, and then click **Add -> Host -> Copy Host**. Grid Manager displays the *Host Record* editor.
  - To modify a host record, select the record and click the Edit icon. Grid Manager displays the *Host Record* editor.
4. The *Host Record* editor provides the following tabs from which you can modify all the fields:
  - **General:** Enter the new hostname and specify at least one IP address. Optionally, you can modify the information you entered through the wizard as described in Adding Host Records above. Note that when you are modifying a DHCP enabled host address, you can apply IPv4 logic filters. To apply IPv4 logic filters, complete the following in the **IPv4 Addresses** section:
    - Select the IP address and click the Edit icon.
    - Click the **Filters** tab in the **Advanced** tab and complete the following:
      - **Filters:** You can keep the inherited IPv4 logic filters or override them and add a new IPv4 logic filter. For information, see [Applying Filters to DHCP Objects](#).
  - **Device Information:** You can change advisory Device Information settings for the object's port reservation; settings are described in the section Adding Host Records above.
  - **TTL:** This tab displays the default TTL settings the record inherited from the Grid or the DNS zone, if you enabled override TTL settings at the zone level. You can keep the default settings or override them. To override the inherited value, click **Override** to enable the configuration. Specify how long the record is cached. Select the time period in seconds, minutes, hours, days, or weeks from the drop-down list. To enable the record to inherit the Grid or zone TTL settings, click **Inherit**.
  - **Aliases:** Click the Add icon. Grid Manager displays a new row in the table. Enter a fully qualified domain name (a CNAME record for the host) in the **Aliases** column. You can delete an alias by selecting the alias checkbox and clicking the Delete icon.
  - **IPv4 Discovered Data:** Displays the discovered data of the IPv4 addresses, if any, of the host record. For information, see [Viewing Discovered Data](#).
  - **Port Reservation:** Review and edit any device port reservations that may be defined for the current object, or create a new port reservation and schedule it. For a closer look, see the section [Port Control Features in Network Insight](#), and steps 4-8 in the section Adding Host Records above.
  - **IPv6 Discovered Data:** Displays the discovered data of the IPv6 addresses, if any, of the host record. For information, see [Viewing Discovered Data](#).
  - **Extensible Attributes:** You can add and delete extensible attributes that are associated with a host record. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [About Administrative Permissions](#).
5. Save the configuration.
6. Click **Restart Services** on the Toolbar for the changes to take effect.

## About Network Containers

Grid Manager uses network containers to group IPv4 and IPv6 networks. A network container is a parent network that contains other network containers and leaf networks. A leaf network is a network that does not contain other networks. For example, the following figure illustrates the IPv4 20.0.0.0/8 network, which is a network container with two network containers, 20.8.0.0/13 and 20.72.0.0/13. The 20.8.0.0/13 network has two leaf networks, 20.8.0.0/16 and 20.9.0.0/16. The 20.72.0.0/13 network has one leaf network, 20.72.0.0/16.

### IPv4 Network Container



From Grid Manager, you can click the link of the network container 20.0.0.0/8 in the IP List panel and drill down to the two network containers, 20.8.0.0/13 and 20.7.0.0/13, as shown in the below figure. You can click the network container links to drill down further to the leaf networks.

### IP List View of Network Containers

NETWORK	CLOUD USAGE	OWNED BY	DELEGATED TO	COMMENT	IPAM UTILIZATION	DISCOVER NOW	DISCOVERY ENGINE
20.8.0.0/13	Non-cloud	Grid			0.0%		None
20.72.0.0/13	Non-cloud	Grid			0.0%		None

In the **IPAM** tab, when you create an IPv4 or IPv6 network that belongs to a larger network, the appliance automatically creates a network container and puts the leaf network in the container. The appliance also creates network containers when you split IPv4 or IPv6 networks into smaller networks. For information, see [Splitting IPv4 Networks into Subnets](#) and [Splitting IPv6 Networks into Subnets](#).

## Adding IPv4 and IPv6 Network Containers and Networks

To add an IPv4 or IPv6 network container or network:

1. From the **Data Management** tab, select the **IPAM** tab.
2. Click the Add icon and select either **IPv4 Network** or **IPv6 Network**.
3. In the *Add Network* wizard, create a network as described in [Adding IPv4 Networks](#) or [Adding IPv6 Networks](#).

## Modifying IPv4 and IPv6 Network Containers and Networks

You can modify existing network settings, with the exception of the network address and subnet mask. To modify an IPv4 or IPv6 network container or network:

1. From the **Data Management** tab, select the **IPAM** tab -> *network* checkbox, and then click the Edit icon.
2. In the *DHCP Network* editor, modify the network settings as described in [Modifying IPv4 Networks](#) or [Modifying IPv6 Networks](#).

## Deleting Network Containers

Depending on the configuration, you may or may not be able to delete or schedule the deletion of a network container and all its contents. Contents in a network container can include other network containers, leaf networks, and associated objects. For recursive deletions, only network containers and networks are considered. Objects such as hosts are not considered for recursive deletions.

Superusers can determine which group of users are allowed to delete or schedule the deletion of a network container and all its contents. For information about how to configure the recursive deletion of network containers, see [Configuring Recursive Deletions of Networks and Zones](#).

Note that you must have Read/Write permission to all the contents in order to delete a network container. When you delete a network container only, the appliance re-parents the other network containers and leaf networks.

The appliance puts all deleted objects in the Recycle Bin, if enabled. You can restore the objects if necessary. When you restore a parent object from the Recycle Bin, all its contents, if any, are re-parented to the restored parent object. For information about the Recycle Bin, see [Using the Recycle Bin](#).

To delete a network container:

1. From the **Data Management** tab, select the **IPAM** tab -> *network\_container* checkbox. You can select multiple network containers for deletion.
2. Click the Delete icon.
3. Do one of the following in the *Delete Confirmation* dialog box:
  - Select one of the following. Note that these options appear only if you are allowed to delete the network container and all its contents. For information about how to configure this, [Configuring Recursive Deletions of Networks and Zones](#).
    - **Delete only the network container and re-parent the subnets:** Select this to delete only the network container and re-parent its subnets.
    - **Delete the network container and all its subnetworks:** Select this to delete both the network and its contents.
  - Click **Yes**.

The appliance puts the deleted network container in the Recycle Bin, if enabled. You can also schedule the deletion for a later time. Click **Schedule Deletion** and in the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Deletions](#). For information about scheduling recursive deletions of network containers, see [Scheduling Recursive Deletions of Network Containers and Zones](#).

## Managing IPv4 Networks

In Grid Manager, you use the Net Map (network map) and List panels to manage your IPv4 network infrastructure. After you select a network container from the IPAM tab, Grid Manager displays it in the Net Map panel, by default. The Net

Map panel provides a graphical view of your networks and has a number of features that simplify network management. The List panel displays the networks in table format.

You can always switch your view of a network container between the Net Map and List panels. Grid Manager keeps track of which panel you last used. When you select a network container, Grid Manager displays it in the Net Map or List panel, depending on which one you last used. For information about each panel, see IPv4 Network Map and IPAM Home below.

Use the IP Map and List panels to manage the IP addresses in leaf networks. For information, see [Viewing and Managing IPv4 Addresses](#).

After you create an IPv4 network, you can modify its properties, resize it, use the split network feature to create subnets, enable discovery to discover routers, switches, firewalls, wireless access points and other device types within it, or join it to another network to create a larger network that encompasses adjacent subnets.

You can do the following from both the Net Map and List panels:

- Resize a network. For information, see [Resizing IPv4 Networks](#) below.
- Split a network into subnets. For information, see [Splitting IPv4 Networks into Subnets](#) below.
- Join a network. For information, see [Joining IPv4 Networks](#) below.
- Discover devices in the network. For information, see [Discovering Networks \(Under Network Insight only\)](#) below.

## IPv4 Network Map

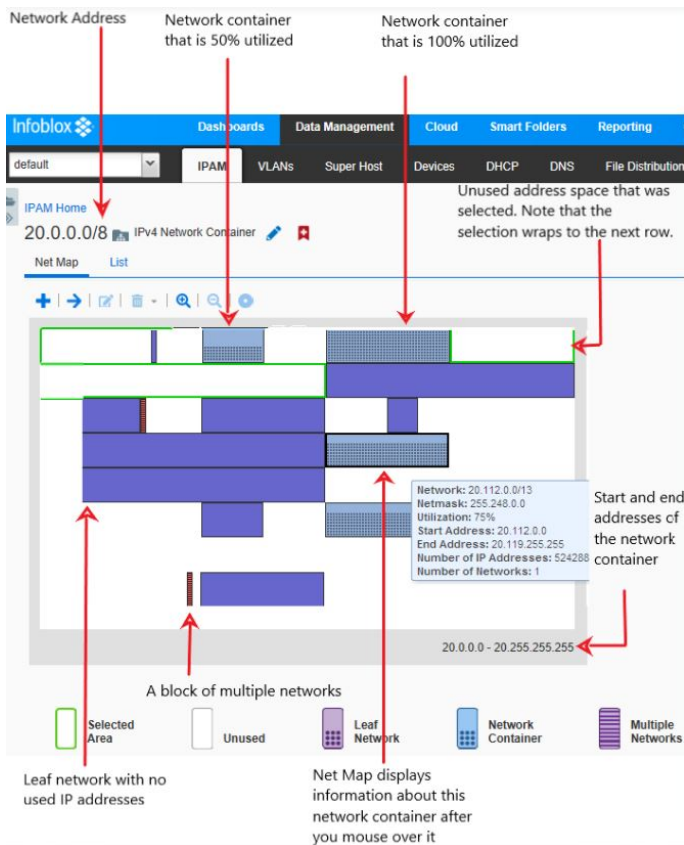
After you select an IPv4 network container from the IPAM tab, Grid Manager displays it in the Net Map (network map) panel, by default. Net Map provides a high-level view of your IPv4 network address space. You can use Net Map to design and plan your network infrastructure, configure and manage individual networks, and evaluate their utilization. Its unique display of the IPv4 network address space across multiple rows is similar to a road map that starts with the first IP address in the network and ends with the last address. Net Map displays the network address space across a maximum of eight rows, depending on the size of the network. It automatically scales the map so that it displays the entire address space of a network container.

The Net Map panel presents a complete view of the network space, including the different types of networks that are in it and its unused address space. IP addresses that belong to a network are blocked off. Each color-coded block represents a network container, a leaf network, or a block of networks that are too small to be displayed individually in the map. For example, in a /8 or /16 network, networks smaller than /20 or /28 respectively and that are beside each other are represented as a multiple network block. In addition, the fill pattern of the blocks indicates their utilization. Therefore, you can quickly evaluate how many and what type of networks are in a network container, their relative sizes, utilization, and how much space you have left.

As you mouse over areas of the map, it displays IP information about the area. Net Map also has a zoom feature that allows you to enlarge or reduce your view of a particular area.

The below figure displays the network map of a 20.0.0.0/8 network, which is a network container that has network containers and leaf networks.

### *20.0.0.0/8 Network Map*



## Displaying IP Information

As shown in 20.0.0.0/8 Network Map figure, as you mouse over the map, Net Map displays IP information about the area. When you mouse over an unused area, Net Map displays the following information:

- The start and end IP address
- The number of IP addresses that can fit in that space
- The largest possible network
- The number of /16 and /24 networks that can fit in that space

When you mouse over a network, Net Map displays the following information:

- Network address and netmask
- Utilization of the network. For a leaf network, Net Map reports the percentage of used IP addresses, except the broadcast and network addresses. For a network container, Net Map reports the percentage of the IP address space that has been allocated to either network containers or leaf networks.
- The first and last IP address of the network
- The total number of IP addresses in the network

When you mouse over a block of multiple networks, Net Map displays the following information:

- The start and end IP address of that block of networks
- The total number of IP addresses in that block of networks
- The number of networks in that block

## Zooming In and Out

Use the zoom function to enlarge and reduce your view of a selected area. You can zoom in on any area in your network. You can zoom in on an area until it displays 128 addresses per row, for a total of 1024 addresses for the map. When you reach the last possible zoom level, the Zoom In icon in the Net Map task bar and the menu item are disabled.

After you zoom in on an area, you can click the Zoom Controller icon to track where you zoomed in. The Zoom Controller lists all the areas that you zoomed in and updates its list dynamically. You can click an item on the list to view that area again. Click the Zoom Controller again to close it.

To select an area and zoom in:

1. Right-click and select **Zoom In**, or click the Zoom In icon in the Net Map task bar. The pointer changes to the zoom in selector.
2. Select a starting point and drag to the end point. The starting point can be anywhere in the map. It does not have to be at the beginning of a network.  
Net Map displays a magnified view of the selected area after you release the mouse button. As you mouse over the zoomed in area, Net Map displays IP information about it.
3. You can do the following:
  - Select an area and zoom in again.
  - Add a network. If you zoom in on an area and click Add without selecting an open area first, Net Map selects the area where it can create the biggest possible network in that magnified area.
  - Select a network and perform any of the following operations:
    - Split the network.
    - Join it to another network.
    - Resize the network.
    - Edit its properties.
    - Open it to display its network or IP map.
  - Right-click and select **Zoom Out**, or click the Zoom Out icon in the Net Map task bar. Each time you click **Zoom Out**, Net Map zooms out one level and the Zoom Controller is updated accordingly.

## Net Map Tasks

From Net Map, you can create IPv4 networks, and evaluate and manage your network resources according to the needs of your organization.

You can do the following:

- Zoom in on specific areas, as described in [Zooming In and Out](#) above.
  - Add a network, as described in [Adding a Network from Net Map](#) below.
  - Select a network and view either its network or Net Map, as described in [Viewing Network Details](#) below.
  - Select a network and edit its properties, as described in [Modifying IPv4 and IPv6 Network Containers and Networks](#).
  - Split a network, as described in [Splitting IPv4 Networks into Subnets](#) below.
  - Join networks, as described in [Joining IPv4 Networks](#) below.
  - Resize a network, as described in [Resizing IPv4 Networks](#) below.
  - (Applies only with *Network Insight*) Execute **vDiscovery** on the selected network, as described in [Configuring vDiscovery Jobs](#).
  - (Applies only with *Network Insight*) View **Discovery Status** for the selected network, as described in [Viewing Discovery Status](#).
  - (Applies only with *Network Insight*) Execute **Discovery Diagnostics** on the selected network, as described in [Executing Discovery Diagnostics](#).
  - (Applies only with *Network Insight*) Direct NIOS to discover devices on the selected network (**Discover Now**). The network must have discovery enabled before this button will be active. For more information about requirements and discovery features, see the topics under [About Network Insight](#).  
Note If the **Discover Now** button and other associated discovery elements are disabled on the Toolbar, it indicates that discovery is not enabled for the parent network of the selected element or IP, or the network is not associated with a discovery appliance.
- Delete one or multiple networks, as described in [Discovering Networks \(Under Network Insight only\)](#) below.



- **Clear All Unmanaged Data** or **Clear All Discovered Data**, as described in the section [Clearing Discovered Data](#).
- Switch to the List view of the network. For information, see [IPAM Home](#) below.
  - When you select one or more networks in Net Map and then switch to the List view, the list displays the page with the first selected network.
  - If you select one or more networks in the List view and then switch to the Net Map view, the first network is also selected in Net Map. If you select a network in the List view that is part of a Multiple Networks block in Net Map, it is not selected when you switch to the Net Map view.

### Adding a Network from Net Map

When you create networks from Net Map, you can evaluate your network infrastructure and add networks accordingly. You can view the address space to which you are adding a network, so you can determine how much space is available and which IP addresses are not in use. When you mouse over an open area, Net Map displays useful information, such as the largest possible network that fits in that area and the total number of IP addresses. In addition, you can create networks without having to calculate anything. When you add a network, Net Map displays a netmask slider so you can determine the appropriate netmask for the size of the network that you need. As you move the slider, it displays network information, including the total number of addresses. After you select the netmask, you can even move the new network around the open area to select another valid start address.

To add a network from the Net Map panel:

1. Do one of the following:
  - Click the Add icon.  
Net Map displays the netmask slider and outlines the open area that can accommodate the largest network.
  - Select an open area, and then click the Add icon.  
Net Map displays the netmask slider and outlines the largest network that you can create in the open area that you selected.
2. Move the slider to the desired netmask. You can move the slider to the netmask of the largest network that can be created in the open area.  
As you move the slider, Net Map displays the netmask and its corresponding number of IP addresses. The outline in the network map also adjusts as you move the slider. When you mouse over the outline, it displays the start and end address of the network.
3. After you set the slider to the desired netmask, you can drag the new network block around the open area to select a new valid starting address. You cannot move the block to a starting address that is invalid.
4. Click **Launch Wizard** to create the network.  
The *Add Network* wizard displays the selected network address and netmask.
5. You can add comments, automatically create reverse mapping zones, and edit the extensible attributes. For information, see [Adding IPv4 Networks](#). You cannot change the network address and netmask, but you can edit the description and enable or disable a network by selecting the network and clicking the Edit icon. To disable a network, you can double click the respective row, select the checkbox in the **Disabled** column and click **Save**. Grid Manager displays a warning message when you select the checkbox. Click **Yes** to confirm or **No** to cancel. You can also delete or restore a network. Grid Manager displays a warning message during deletion and when you restore the network indicating that the process may take a longer time if the amount of data is huge. Click **Yes** to continue or **No** to cancel the process.
6. Save the configuration and click **Restart** if it appears at the top of the screen. Grid Manager updates Net Map with the newly created network.

### Viewing Network Details

From the Net Map panel, you can focus on a specific network or area and view additional information about it. If you have a network hierarchy of networks within network containers, you can drill down to individual leaf networks and view their IP address usage.

1. Select a network or area.
2. Click the Open icon.
  - If you selected a network container, Grid Manager displays it in the Net Map panel. You can drill down further by selecting a network or open area and clicking the Open icon again.

- If you selected a block of multiple networks, Grid Manager displays the individual networks in the Net Map panel. You can then select a network or open area for viewing.
- If you selected a leaf network, Grid Manager displays it in the IP Map panel.
- If you selected an open area, Grid Manager displays an enlarged view of that area in the Net Map panel.

This is useful when you are creating small networks in an open area.

## IPAM Home

The IPAM Home panel is an alternative view of an IPv4 and IPv6 network hierarchy. For a given network, the panel shows all the networks of a selected network view in table format. This panel displays only the first-level subnets. It does not show further descendant or child subnets. You can open a subnet to view its child subnets. Subnets that contain child subnets are displayed as network containers. If the number of subnets in a network exceeds the maximum page size of the table, the network list displays the subnets on multiple pages. You can use the page navigation buttons at the bottom of the table to navigate through the pages of subnets.

The IPAM home panel displays the following:

- **Network:** The network address.
- **Comment:** The information you entered about the network.
- **RIR Organization:** This appears only if support for RIR updates is enabled. This displays the name of the RIR organization to which the network is assigned.
- **RIR OrganizationID:** This appears only if support for RIR updates is enabled. This displays the ID of the RIR organization to which the network is assigned.
- **RIR RegistrationStatus:** This appears only if support for RIR update is enabled. This field displays the RIR registration status. This can be **Registered** or **Not Registered**. **Registered** indicates that the network has a corresponding entry in the RIPE database.
- **Last Registration Updated:** Displays the timestamp when the last registration was updated. The displayed timestamp reflects the timestamp used on the Grid Master.
- **Status of Last Registration Update:** Displays the registration status and communication method of the last registration update. The status can be Pending, Sent, Succeeded, or Failed. Each time you send a registration update to create, modify, or delete a network container or network, the updated status will be displayed here. If you have selected not to send registration updates, the previous status is retained.
- **IPAM Utilization:** For a network, this is the percentage based on the IP addresses in use divided by the total addresses in the network. For example, in a /24 network, if there are 25 static IP addresses defined and a DHCP range that includes 100 addresses, the total number of IP addresses in use is 125. Of the possible 256 addresses in the network, the IPAM utilization is about 50% for this network.  
For a network container that contains subnets, this is the percentage of the total address space defined within the container regardless of whether any of the IP addresses in the subnets are in use. For example, when you define a /16 network and then 64 /24 networks underneath it, the /16 network container is considered 25% utilized even when none of the IP addresses in the /24 networks is in use.  
You can use this information to verify if there is a sufficient number of available addresses in a network. The appliance updates the IPAM utilization data immediately for a network container, but for a network it is updated every 15 minutes.  
The IPAM utilization data is displayed in one of the following colors:
  - Red: The IPAM utilization percentage is above the configured Trigger value.
  - Blue: The IPAM utilization percentage is below the configured Trigger value.
- **Active Users:** The number of active users on the selected network.
- **Discovery Engine:** Displays the discovery engine that performs the discovery process. This can be **Network Insight**, **NetMRI**, or **vDiscovery**. This field displays **None** if you have added or imported the network container or network manually.
- **Site:** The site to which the IP address belongs. This is a predefined extensible attribute.
- **Bridge Domain:** The name of the discovered bridge domain. This column will display values only for IP addresses that are discovered from Cisco APIC by Network Insight or NetMRI. If discovered by NetMRI, the value will be populated through IPAM Sync. Otherwise, this column will be blank. For information about how to configure Cisco APIC, see [Configuring Cisco Application Policy Infrastructure Controller \(APIC\)](#).
- **Tenant:** The name of the discovered tenant. This column will display values only for IP addresses that are discovered from Cisco APIC by Network Insight or NetMRI. If discovered by NetMRI, the value will be populated



through IPAM Sync. Otherwise, this column will be blank. For information about how to configure Cisco APIC, see [Configuring Cisco Application Policy Infrastructure Controller \(APIC\)](#).

Additionally, you can select the following columns for display:

- **Disabled:** Indicates whether the network is disabled.
- **Leaf Network:** Indicates whether the network is a leaf network or not. A leaf network is a network that does not contain other networks.
- **Discovery Enabled:** (*Applies only with Network Insight*) Indicates whether discovery is allowed on the network container or the network.
- **Discover Now:** (*Applies only with Network Insight*) Indicates when the network is undergoing a current discovery process. A "Pending" icon appears in this column when you start the discovery and displays **Completed** after the completion of the discovery process.
- **Discovered VLAN ID:** (*Applies only with Network Insight*) The VLAN ID on the switch port.
- **Discovered VLAN Name:** (*Applies only with Network Insight*) The VLAN name on the switch port.
- **Assigned VLAN ID:** VLAN ID of the VLAN object assigned to the network.
- **Assigned VLAN Name:** VLAN name of the VLAN object assigned to the network.
- **VRF Name:** The name of the discovered VRF that uses IP addresses of the network, as discovered by Network Insight or NetMRI.
- **VRF Description:** The description of the discovered VRF.
- **VRF RD:** The address of the route distinguisher of discovered VRF.  
Note that as a general rule for the VRF-related columns, a column displays a specific value if there is a single non-empty value or several same values for the IP addresses in the network. Otherwise, the column displays "Multiple". For example, if the VRF names for a network have the same value, the **VRF Name** column displays this value for the network. If more than one VRF are discovered for a network and their names are different, the **VRF Name** column displays "Multiple". However, if for a number of VRFs there is only one VRF description or VRF RD value among other empty strings, the columns **VRF Description** and **VRF RD** display "Multiple" as this is regarded as distinct VRFs.  
To see values for each IP address in the network, click the network -> **List** tab.
- **BGP AS:** The number of the discovered BGP Autonomous System that uses IP addresses of the network. Note If more than one BGP AS are discovered for a network and their numbers are different, the **BGP AS** column displays "Multiple". If the BGP AS numbers have the same value, the **BGP AS** column displays this value for the network. To see values for each IP address in the network, click the network -> **List** tab.
- **Managed:** (*Applies only with Network Insight*) Indicates whether the network is set to Managed status under Grid Manager. For more information, see the section [Converting Unmanaged Networks under IPAM to Managed Status](#).
- **First Discovered:** (*Applies only with Network Insight*) The date and timestamp of the first occasion that Grid Manager discovered the network.
- **Last Discovered:** (*Applies only with Network Insight*) The date and timestamp of the last occasion that Grid Manager performed discovery on the network. The timestamp is updated whenever any new IP from this network is discovered.
- **Extensible attributes and RIR attributes:** You can select the extensible attributes such as Building, Country, Region, State, and VLAN for display. When you enable support for RIR registration updates, you can also select associated RIR attributes for display. For information about RIR attributes, see [Managing RIR Attributes](#).
- **Active Directory Sites:** You can also select Active Directory Sites for display. For information about Active Directory Sites, see [About Active Directory Sites and Services](#).

The Toolbar and the Action menu provide access to actions you can take on the selected network where applicable, as follows:

- Click **Show Device View** to view the devices whose interfaces are in the selected network. To view device details and a list of all the interfaces associated with it, click the name of the device.
- Click **Convert** to change the status of an unmanaged network to managed status under IPAM.
- Click **Discover Now** to apply Network Insight discovery to a listed network.
- Click **Edit** to open the network editor.
- Click **Delete** to delete or schedule deletion of the selected network. To delete the network now, in the Delete Confirmation dialog box, click Yes. Grid Manager displays a warning message. Click Yes to continue or No to cancel the process.
- Click **Extensible Attributes** to open the network editor's Extensible Attributes page for the selected network.
- Click **Permissions** to open the network editor's Permissions page for the selected network.

- Click **Move Networks** to move a network to a destination Active Directory site.
- Click **Show Active Users** to view all the users who are currently active on the Active Directory domain.

You can sort the list of networks in ascending or descending order by columns. For information about customizing tables in Grid Manager, see [Customizing Tables](#).

You can also modify some of the data in the table fields. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).

---

**Tip:** If you select a network from the list and switch to the Net Map panel, the network is also selected in the network map.

---

### Filtering the Network List

You can filter the network list, so it displays only the networks you need. You can filter the list based on certain parameters, such as network addresses, comments and extensible attributes. When you expand the list of available fields you can use for the filter, note that the extensible attributes are those with a gray background.

### Resizing IPv4 Networks

You can resize a network to increase or decrease the network size and the number of IP addresses in the network. When you resize a network to a smaller netmask, you increase the number of IP addresses within that network. You can change the size of an IPv4 network when the operation does not affect existing objects in the network. You can resize an existing network only if the resized network does not exceed the upper network limit or create orphan objects, such as hosts and DHCP ranges. When a network has a parent network or subnets, the upper limit of the network size is marked in red in the resize network slider, and you cannot resize beyond this limit. For example, if a network has a /16 parent network, you cannot resize the network to a network that is larger than /16.

Before you resize an RIR allocated network block, ensure that the network block has already been registered at the corresponding RIR. Otherwise, when you reassign addresses within this block, the registration updates may fail. For information about RIR registration updates, see [RIR Registration Updates](#).

To resize a network:

1. From the Net Map or List panel, select a network, and then click **Resize** from the Toolbar.
2. In the *Resize Network* editor, do the following:
  - **Address:** Displays the network address. You cannot modify this field.
  - **Netmask:** Displays the netmask of the network as you resize the network. You cannot modify this field.
  - **Resize slider:** Use the resize network slider to specify the appropriate subnet masks for the subnets. When you move the slider, Grid Manager displays the number of subnets and IP addresses within that subnet.
  - **Automatically create reverse-mapping zone:** This is enabled only when you resize a /8, /16, or /24 network. Select this checkbox to have the appliance automatically create reverse-mapping zones for the subnet. The appliance automatically creates reverse-mapping zones only for /8, /16, and /24 netmasks.
3. Click **OK**.

### Splitting IPv4 Networks into Subnets

You can create smaller subnets simultaneously within a network by splitting it. You do not have to configure each subnet individually. You can create smaller subnets with larger netmasks. A larger netmask defines more networks with a smaller number of IP addresses.

These subnets inherit the address properties of the parent network, such as member assignments. The exceptions are the default router and broadcast address configuration. The default router and broadcast address configuration for address ranges and fixed address are disabled by default after splitting a network. You can enable these properties for each subnet after splitting the parent network.

Note that you cannot split a network that is part of a shared network.

To split a network:

1. From the Net Map or List panel, select the checkbox of a network, and then click **Split** from the Toolbar.
2. In the *Split Network* editor, do the following:
  - **Address:** Displays the network address. You cannot modify this field.
  - **Netmask:** Displays the netmask of the network. You cannot modify this field.
  - **Subnetworks:** Displays the number of subnets and IP addresses for each subnet.
  - **Split network slider:** Use the split network slider to specify the appropriate subnet masks for each subnet. When you move the slider, Grid Manager displays the number of subnets and the IP address range within that subnet.
  - **Immediately Add:** Select one of the following options.
    - **Only networks with ranges and fixed addresses and unmanaged:** Adds only the networks that have DHCP ranges, fixed addresses, and unmanaged addresses.
    - **All possible networks:** Adds all networks that are within the selected netmasks. This is enabled only when you split the /8 networks to /9 or /16 networks. Note that when you add a large number of networks, it could take a little longer for Grid Manager to display the networks.
  - **Automatically create reverse-mapping zone:** Select this checkbox to have the appliance automatically create reverse-mapping zones for the subnets.
3. Click **OK**.

## Joining IPv4 Networks

Joining multiple networks into a larger network is the opposite of splitting a network. You can select a network and expand it into a larger network with a smaller netmask. A smaller netmask defines fewer networks while accommodating a larger number of IP addresses. Joining or expanding a network allows you to consolidate all of the adjacent networks into the expanded network. Adjacent networks are all networks falling under the netmask of the newly-expanded network. You can expand the selected network to a new size and add all other subnets into the new network. When you join networks, you need not define all small networks that cover the address spaces for a larger network.

Each of the adjacent networks join the expanded network and inherit the DHCP member configuration options of the selected network. The expanded network does not inherit the default router and broadcast address configurations of the adjacent networks. Those configurations are disabled by default.



### Note

The member assignment for the expanded network combines all member assignments of the joining networks.

Note that the join and resize features work identically only when you have a single network. If the resize feature is disabled and if you have a single network object with additional new networks, then you must use the join feature to combine all networks.

To join or expand a network:

1. From the Net Map or List panel, select a network, and then click **Join** from the Toolbar.
2. In the *Join Network* editor, do the following:
  - **Address:** Displays the network address. You cannot modify this field.
  - **Netmask:** Displays the netmask of the network as you expand the network.
  - **Join Network slider:** Use the join network slider to specify the available subnet masks for the newly expanded network. Select a smaller netmask value, based on your requirements of the newly-expanded network. When you move the slider, a dialog box displays the total number of IP addresses and the IP address range of a selected subnet mask.

- **Automatically create reverse-mapping zone:** Select this checkbox to configure the expanded network to support reverse-mapping zones [Adding Grid Members](#).
3. Click **OK**.

## Discovering Networks (Under Network Insight only)



### Note

When you add a new network and select the **Enable Discovery** option in the **Add IPv4 Network Wizard** window, network discovery begins automatically, and you do not need to click **Discover Now**.

If the **Discover Now** button and other associated discovery elements are disabled on the Toolbar, it indicates that discovery is not enabled for the parent network of the selected network or IP, or that a discovery appliance (known as a Probe) is not associated with the network that you wish to discover.

To discover IPv4 or IPv6 networks:

- From the Net Map or List panel, select a network, and then click **Discover Now** from the Toolbar. NIOS asks you to confirm that you wish to launch discovery on the selected network.

In the Net Map panel, you can click on IP addresses in the network being discovered. As new data becomes available, NIOS updates the Discovered Data section of the panel with any information found on the device associated with the selected IP.

For more information about requirements and discovery features, see the topics under [About Network Insight](#).

## Deleting Networks

From the IPAM tab, you can delete multiple IPv4 and IPv6 networks. When you delete a network, all of its data, including all of its DHCP records, subnets, and records in its subnets, is deleted from the database and goes to the Recycle Bin, if enabled. Because of the potentially large loss of data that can occur when you delete a network, Grid Manager requires a confirmation to move the data to the Recycle Bin.

To delete IPv4 or IPv6 networks:

1. From the **Data Management** tab, select the **IPAM** tab -> *network* checkbox. You can select multiple checkboxes for multiple networks.
2. Select **Delete** or **Schedule Delete** from the Delete drop-down menu.
3. To delete the network now, in the *Delete Confirmation* dialog box, click **Yes**. To schedule a deletion, see [Managing Extensible Attributes](#).

The appliance puts the deleted network in the Recycle Bin, if enabled.

## Viewing and Managing IPv4 Addresses

You can view and manage IPv4 address data in the IP Map and List panels. Grid Manager displays the IP Map and List panels for a specific network after you navigate through the network hierarchy, or when the selected network does not have subnets under it.

### IP MAP

The IPv4 Map panel provides a graphical representation of all IPv4 addresses in a given subnet. IP Map displays cells that represent IPv4 addresses. Each cell in the map represents an IPv4 address, and its color indicates its status as described in the legend section. You can run a network discovery on the selected network, and the status of each IP address is updated accordingly. For information, see [About Discovery](#).

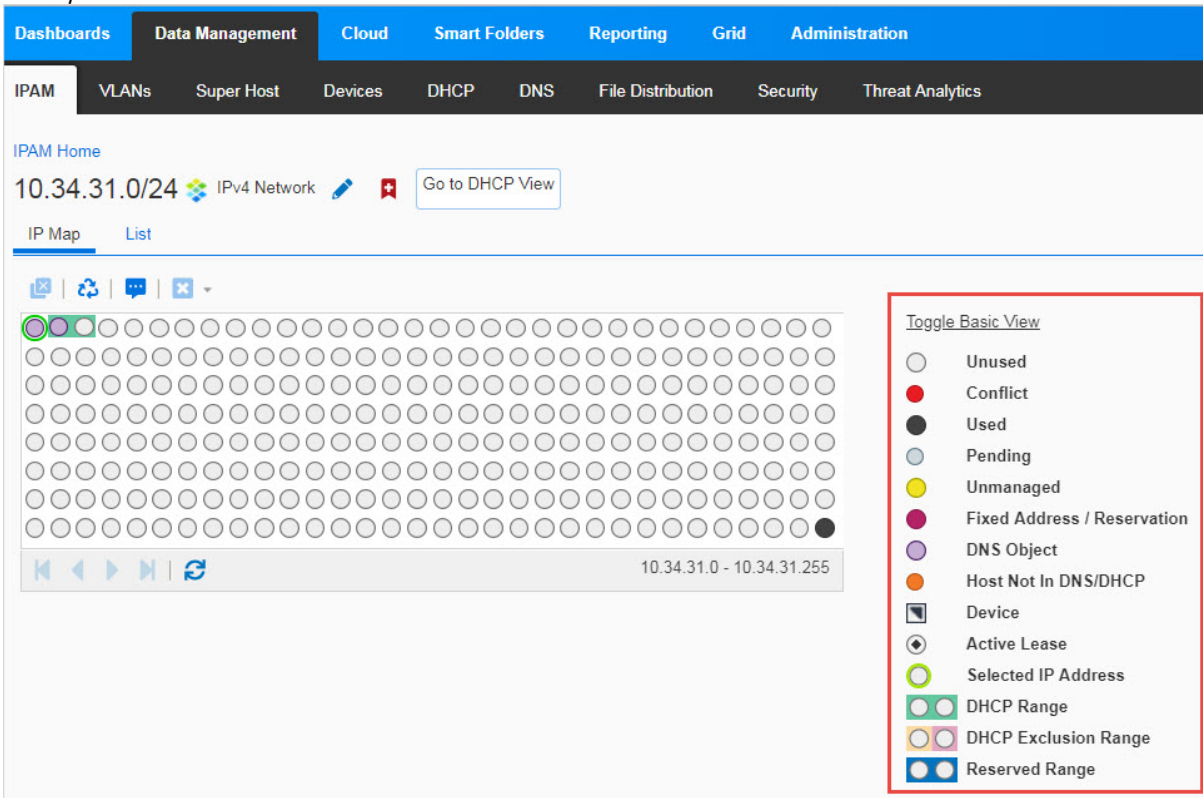
Each IP Map panel can accommodate up to 256 cells with each cell representing an IP address. If a given network has more than 256 addresses, additional IP addresses are displayed by paging to the next page. You can use the page navigation buttons to page through the IP addresses. To go to a specific IP address, you can enter the IP address in

the **Go to** field or click a specific cell in IP Map. IP Map has a basic and an advanced view. You can toggle between these views by clicking **Toggle Basic View** or **Toggle Advanced View**. As illustrated in the figure IP Map - Advanced View below, the status of an IP address is represented with a different color in the IP Map panel.

In the basic view, the IP Map panel displays the following IP address status:

- **Unused:** An IP address that has not been detected and is not associated with any network device or active host on the network.
- **Conflict:** An IP address that has either a MAC address conflict or a DHCP lease conflict detected through a network discovery.
- **Used:** An IP address that is associated with an active host on the network. It can be a resource record, fixed address, reservation, DHCP lease, or host record.
- **Pending:** An IP address that is associated with a scheduled task or approval workflow, and the associated operation has not been executed yet. This IP address is not considered when using the next available IP address function.
- **Selected IP Address:** The IP address that you selected.
- **DHCP Range:** The IP addresses within a DHCP range in the network. The appliance highlights the cells using a green background.
- **Reserved Range:** A range of IP addresses that are reserved for statically configured hosts. They are not served as dynamic addresses. You can allocate the next available IP from the reserved range when you create a static host.

### IP Map - Advanced View



In the advanced view, the IP Map panel displays additional status as follows:

- **Unmanaged:** An IP address that has a discovered host, is not previously known to the appliance, and does not have an A record, a PTR record, fixed address, host address, lease, or is not within a DHCP range. You can change an unmanaged address to a host, DHCP fixed address, A record, or PTR record. You can also clear an unmanaged address. All existing administrator permissions apply to the unmanaged addresses.
- **Fixed Address/Reservation:** A host that is either a fixed address or reservation.
- **DNS Object:** An object that is configured for DNS usage.

- **Host Not in DNS/DHCP:** An IP address that is associated with a host record, but is not configured for DHCP or DNS services.
- **Active Lease:** An IP Address that has an active DHCP lease.
- **DHCP Exclusion Range:** A range of IP addresses within a DHCP range. The appliance cannot assign addresses in the exclusion range to a client. You can use these addresses as static IP addresses. This prevents address conflicts between statically configured devices and dynamically configured devices.  
Note that for a Microsoft split-scope range, the appliance highlights the cells using a combination of orange and pink background colors when the network is managed by two Microsoft servers. For a DHCP exclusion range, the appliance highlights the cells using an orange background.
- **Reserved Exclusion Range:** A range of IP addresses that are reserved for statically configured hosts. IP addresses residing within the exclusion range are excluded from the pool of available IP addresses and are not available for lease.

Under the IP map, Grid Manager displays the following information for the IP address that you have selected in the map:

- **Type:** The object type that is associated with the IP address. For example, this can be **Lease**, **IPv4 DHCP Range** or **Fixed Address**.
- **Comment:** Additional information about the IP address.
- **Lease State:** The lease state of the IP address. This can be one of the following: **Free**, **Backup**, **Active**, **Expired**, **Released**, **Abandoned**, **Reset**, **BootP**, **Static**, **Offered**, or **Declined**.
- **Name:** The name of the object type associated with the IP address. This field displays the name of the object type in the native character set if a host record contains IDNs. If a host record contains IDNs in punycode, this field displays the name in the punycode representation. For example, if the IP address belongs to a host record, this field displays the hostname. For IDNs, this field displays the name in the native character set. If punycode is used, then the appliance displays name in punycode.
- **MAC Address:** The discovered MAC address of the host. This is the unique identifier of a network device. The discovery acquires the MAC address for hosts that are located on the same network as the Grid member that is running the discovery. This can also be the MAC address of a virtual entity on a specified vSphere server. The appliance displays an **X** mark beside the MAC address if it is invalid. For more information about invalid MAC addresses, see [Synchronizing IP Addresses with Invalid MAC Addresses](#).
- **DHCP Fingerprint:** The name of the DHCP fingerprint or vendor ID of the network device that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#).

Depending on the selected object, you may also see information about neighboring devices, port, etc. For more information, see IP List Neighbor Information below.

You can do the following in the IP Map panel:

- Click **Go to DHCP View** to view DHCP properties of a selected network.
- Select an address range by clicking once on a start address and then use SHIFT+click on the end address. Click **Add -> Range** from the Toolbar to add the selected range as an IPv4 or IPv6 DHCP range or reserved range.
- Click the Resolve Conflict icon to resolve IP address conflicts. For information, see [Resolving Conflicting Addresses](#).
- Click the Ping icon to ping a selected IP address. For information, see [Pinging IP Addresses](#).
- Click the Reclaim icon to reclaim an IP address. For information, see [Reclaiming Objects Associated with IPv4 and IPv6 Addresses](#).
- Click the Clear icon to clear an active lease. For information, see [Clearing Active DHCP Leases](#). You can also select an IP address from the IP Map panel and view the following information:
- General information, as described in IP Address Header Panel below.
- Data retrieved through a network discovery or integrated from a PortIQ appliance and Trinzic NetMRI appliance. For information, see [Viewing Discovered Data](#).
- The records associated with the IP address, as described in [Related Objects](#).
- The audit history, as described in [Audit History](#).
- Detailed lease information, as described in [Viewing Detailed Lease Information](#).
- Click **DHCP View** to view DHCP properties of the selected network. For information, see [Modifying IPv4 Networks](#).
- View active network users, as described in [Viewing Active Network Users](#).



## IP Address List

The IP address **List** panel displays all IPv4 addresses of a selected subnet in table format. The list provides information about the IP addresses in a hierarchy view. You can use this list to view detailed information about each IP address and its related objects in a selected network. This list provides information such as address status, object type, and usage.

You can configure filter criteria to display only IP addresses that you want to see in the table. For example, you can enter "MAC Address begins with 00" as the filter criteria to view only IP addresses that have associated MAC addresses that begin with 00. You can also enter a specific IP address in the **Go to** field to view information about the address.

Grid Manager can display the following information for the IP addresses. You can edit the columns to display information that is not shown by default.

- **IP Address:** The IP address of the corresponding record. The appliance highlights disabled DHCP objects in gray. A DHCP object can be an DHCP address range, fixed address, reservation, host configured for DHCP, or roaming host with an allocated IP address.
- **Name:** The name of the object type associated with the IP address. For example, if the IP address belongs to a host record, this field displays the hostname.
- **MAC Address:** The discovered MAC address of the host. This is the unique identifier of a network device. The discovery acquires the MAC address for hosts that are located on the same network as the Grid member that is running the discovery. This can also be the MAC address of a virtual entity on a specified vSphere server.
- **Bridge Domain:** The name of the discovered bridge domain. This column will display values only for IP addresses that are discovered from Cisco APIC through Network Insight or NetMRI. If discovered by NetMRI, the value will be populated through IPAM Sync. Otherwise, this column will be blank. For information about how to configure Cisco APIC, see [Configuring Cisco Application Policy Infrastructure Controller \(APIC\)](#).
- **Tenant:** The name of the discovered tenant. This column will display values only for IP addresses that are discovered from Cisco APIC through Network Insight or NetMRI. If discovered by NetMRI, the value will be populated through IPAM Sync. Otherwise, this column will be blank. For information about how to configure Cisco APIC, see [Configuring Cisco Application Policy Infrastructure Controller \(APIC\)](#).
- **EPG:** The name of the end point group. This column will display values only for IP addresses that are discovered from Cisco APIC through Network Insight or NetMRI. If discovered by NetMRI, the value will be populated through IPAM Sync. Otherwise, this column will be blank. For information about how to configure Cisco APIC, see [Configuring Cisco Application Policy Infrastructure Controller \(APIC\)](#).
- **DHCP Client Identifier:** For an IPv4 address, the DHCP Unique Identifier of the host.
- **Port Reservation:** Lists any Port Reservation from Network Insight that is associated with the IP address. The information takes the form of *device name:interface name*.
- **VIP:** Indicates when the IP address is operating as a Virtual IP and operates in router redundancy.
- **Status:** The current status of the corresponding record. This can be **Used** or **Unused**.
- **Type:** The object type that is associated with the IP address. For example, this can be **Broadcast**, **Lease**, **IPv4 DHCP Range** or **Fixed Address**.
- **Discover Now:** Indicates when the network is undergoing a current discovery process. A "Pending" icon appears in this column when you start the discovery and displays **Completed** after the completion of the discovery process.
- **Usage:** Indicates whether the IP address is configured for DNS or DHCP.
- **Lease State:** The lease state of the IP address. This can be one of the following: **Free**, **Backup**, **Active**, **Expired**, **Released**, **Abandoned**, **Reset**, **BootP**, **Static**, **Offered**, or **Declined**.
- **User Name:** The name of the user who received the lease for the IP address.
- **Task Name:** The name of the task that collected the discovered data. It is usually the ID or task name that collected the data. It is defined on the corresponding Trinzic NetMRI appliance when you import the discovered data to the NIOS appliance. The task name should be defined in the vDiscovery task manager for vDiscovery.
- **Comment:** Additional information about the IP address.
- **First Discovered:** The timestamp when the IP address was initially discovered. This data is read-only.
- **Last Discovered:** The timestamp when the IP address was last discovered. This data is read-only.
- **OS:** The operating system of the discovered host. Sometimes this field also displays the percentage of certainty about the discovered OS. The OS value can be one of the following:
  - **Microsoft:** This value is displayed for all discovered hosts that have a non-null value in the MAC addresses using the NetBIOS discovery method.
  - A value that a TCP discovery returns.
  - The OS of a virtual entity on a vSphere server.

- DHCP fingerprint.
- Information about the device OS from Network Insight. For network devices, OS information is collected by SNMP. Depending on the device SNMP settings, this field can display the OS version or remain empty (rarely).

For end hosts, the OS information is collected using active TCP fingerprint scanning by nmap. To enable this, select the following checkboxes in the Grid basic discovery properties: **Port Scanning** and **Profile Device**.

For Windows end hosts, you can enable Network Insight to use the SMB v1 protocol during fingerprinting. To enable this, go to **Grid Discovery Properties** -> **Advanced** and add port 445 with smb1 service and TCP type.

If nmap fails to detect the OS precisely, a list of most likely guesses is displayed. The end host's **Device Type** field always displays the OS family.

- **NetBIOS Name:** The returned NetBIOS name from the last discovery.
- **Device Type(s):** The type of device associated with the IP address, if any: Router, Switch-Router, and other types.
- **Open Port(s):** Lists any TCP/UDP ports that are open on the current IP address.
- **VRF Name:** The name of the discovered VRF to which the interface with this IP address belongs, as discovered by Network Insight or NetMRI.
- **VRF Description:** The description of discovered VRF.
- **VRF RD:** The address of the route distinguisher of discovered VRF.
- **BGP AS:** The number of the discovered BGP Autonomous System that uses the IP address.
- **Discovered Name:** The name of the discovered IP address, if any was previously assigned by an administrator.
- **Discoverer:** The identity of the appliance that discovered the IP address.
- **Fingerprint:** The name of the DHCP fingerprint or vendor ID of the network device that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#).
- **Site:** The site to which the IP address belongs. This is a predefined extensible attribute.

**Disabled** (hidden): Indicates whether the DHCP or DNS record is disabled.



#### Note

For an IP address that falls within a DHCP range, Grid Manager displays extensible attribute values for the DHCP range and fixed address or host record. When you view the same IP address in the **DHCP** tab however, Grid Manager displays only the extensible attribute values associated with the fixed address or host record, but not the DHCP range. For example, when you define extensible attribute **State** with the value **California** for DHCP range 1.0.0.1 – 1.0.0.5, and then define extensible attribute **State** with the value of **Alaska** for fixed address 1.0.0.3, Grid Manager displays both **California** and **Alaska** in the **State** field for IP address 1.0.0.3 in the IP Address List view. However, when you view 1.0.0.3 from the **DHCP** tab, the **State** field displays **Alaska** only.

## IP List Neighbor Information

Other values are available for display in the IP list. Most of the hidden columns consist of information discovered by Grid Manager, to identify devices connected to the network interfaces neighboring the listed IP addresses. These data columns are hidden and must be enabled for display by selecting the **Visible** checkbox for each field when editing the columns for the table. Most of this information does not appear unless a device is recognized in IPAM or is in managed state under IPAM. Additional IP list neighbor information includes:

- **Discovered Name:** The discovered name of the device bound to the IP address. The IP address may be just one of several or many IP addresses bound to the device on a specific interface.
- **Discovered MAC Address:** The discovered MAC of the interface bound to the IP address.
- **Discoverer:** Name of the Infoblox appliance that discovered the IP address and its associated information.
- **Attached Device Description:** Listing of the device neighboring to the IP address.
- **Attached Device Address:** The IP address of the device neighboring to the current IP.
- **Attached Device Name:** The host name of the device neighboring to the current IP.
- **Attached Device Vendor:** The vendor of the device neighboring to the current IP address.
- **Attached Device Model:** The vendor's model number for the device neighboring to the current IP address.



- **Attached Device Port Description:** An admin-provided description for the neighboring IP address, if any is discovered.
- **Attached Device Port Name:** The name of the switch port or port channel connected to the end device. For Cisco devices with virtual port channel configured, this field also displays the list of physical interfaces that form the virtual port channel. The first vPC peer's IP address is displayed in the field **Attached Device Address** mentioned above, and other vPC data is displayed in the **Attached Device Port Name** field in the following format:  
*virtual port channel name: list of physical interfaces of the first device, IP address of the second device/virtual port channel name: list of physical interfaces of the second device*


Example:

An end device is connected to devices 192.168.1.2 and 192.168.1.3 using vPC. The first vPC peer's IP address 192.168.1.2 is displayed in the **Attached Device Address** field. Other vPC data is displayed in the **Attached Device Port Name** field:

```
port-channel-1: Ethernet4/11, Ethernet5/11, 192.168.1.3/port-channel-1:
Ethernet3/12, Ethernet4/12
```


- **Attached Device Type:** Indicates the device type for the neighboring device: **Router**, **Switch-Router**, **Switch**, and other types.
- **Port Duplex:** Discovered Duplex setting for the neighboring port, when applicable.
- **Port Link:** Indicates the state of the link: **Connected** or **Disconnected**.
- **Port Speed:** Indicates the speed of the network connection.

## Using the List Panel Action Menu

The IP address **List** panel provides an Action icon  with a series of menu options for features related to IP address management and IP data management under IPAM. Menu choices change based upon the context and the current state of IP addresses in the table; features available in the List panel action menu include the following:

- (Applies only with Network Insight) View **Router Redundancy** information for discovered IP addresses in an IPv4 network. For active VIPs, you will see several sets of related information:

### Virtual IPs and discovered redundancy information



- **Active:** lists the active interface in the redundancy pair;
- **VIP:** The Virtual IP for the router redundancy pair;
- **Standby:** The standby IP interface for the router redundancy.

Discovery of all three IP components of the Router Redundancy instance also provides related information for all three IP entities:

- **Show IPAM Address:** opens the IPAM page to the listed IP address;
- **VIP:** Opens the virtual interface in the host device's Interfaces page;
- **Associated router:** The third item lists the hostname router for each of the three IP address entities comprising the redundancy instance. The currently active router will be identified with the **Active** and **VIP** objects; the second **Standby** router is identified with the **Standby** IP address.
- **Show Interfaces View:** Displayed only if there is a device interface associated with this IP address. It displays the interface name. Clicking the interface name opens the interfaces list view for that device. The interfaces list view is pre-populated with the interface name.

- **Assets View → Show Assets:** Opens the IP address's list of network assets that are connected to the IP address or reachable through the IP address in some way, such as through a routable path. Provides a quick look at basic connectivity provided by the selected IP address; for more information about Assets views, see [Viewing Assets Associated with Discovered Devices](#).
- **Clear:** Allows you to remove data associated with the currently selected IP address, from three categories:
  - **Clear Lease:** If the IP interface gets its configuration from a DHCP lease, choosing this option will clear the IP's DHCP configuration. Applies only to IP addresses that are fully managed through IPAM;
  - **Clear Unmanaged Data:** Clears the discovered data for an unmanaged IP address;
  - **Clear Discovered Data:** Clears discovered data from the IPAM object, and re-launches discovery afterwards if necessary.
- **Convert:** Conversion feature to convert the currently unmanaged IP address to an object fully managed by IPAM: **Network, To Host, To A Record, To PTR Record, or To Fixed Address**. (For related information, see [Managing Discovered Data, About Host Records, Managing A Records](#) and [Managing PTR Records](#).)
- **Discover Now:** Requests Grid Manager to execute discovery on the selected IP address.
- **Exclusion:** Exclude or disable exclusion of the current IP address from discovery. For related information, see [Excluding IP Addresses from Discovery](#).
- **Show Active Users:** Displays all the users who are currently active on the selected network. For information, see [Viewing Active Network Users](#).

### Additional IP List Information

You can display all available extensible attributes. You can also sort the list of IP addresses in ascending or descending order by **IP Address** only. If you enabled the IP Discovery feature, you can configure the IP List panel to display discovered data and fields imported from NetMRI appliances. For information about integrating discovered data from NetMRI, see [Integrating Discovered Data From NetMRI](#).

You can select an IP address from the List panel and view the following information about it:

- General information, as described in [IP Address Header Panel](#).
- Data retrieved through a network discovery or integrated from a PortIQ appliance, as described in [Viewing Discovered Data](#).
- The records associated with the IP address, as described in [Related Objects](#).
- Audit history, as described in [Audit History](#).
- Detailed lease information, as described in [Viewing Detailed Lease Information](#).

You can also do the following from the IP List panel:

- Click **Go to DHCP View** to view DHCP properties of a selected network. For information, see [Modifying IPv4 Networks](#).
- Click the Ping icon to ping a selected IP address. For information, see [Pinging IP Addresses](#).

### Filtering the IP Address List

You can filter the IP address list, so it displays only the IP addresses you need. You can filter the list based on any combination of extensible attributes and the parameters displayed in the IP address list, such as usage and type. When you expand the list of available fields you can add to the filter, note that the extensible attributes are those with the gray background.

### IP Address Header Panel

When you select an IP address from the IP Map or List panel, Grid Manager displays information about the highest priority object associated with the IP address. Depending on the object type, Grid Manager displays all or some of the following information. For example, if the highest priority object is a fixed address, Grid Manager displays only the object type, MAC address, lease state, and comment of the object.

- **Type:** The object or record type, such as A record, PTR record, or host record.

- **Name:** The name of the object. For example, if the IP address belongs to a host record, this field displays the hostname. The appliance highlights disabled DHCP objects in gray. A DHCP object can be a DHCP address range, fixed address, reservation, host configured for DHCP, or roaming host with an allocated IP address.
- **MAC:** The MAC address of the network device associated with the IP address.
- **Lease State:** The current status of the DHCP lease.
- **Comment:** Comments about the IP address.

## Discovered Data

The **Discovered Data** tab displays discovered data through a network discovery or integrated from PortIQ and NetMRI appliances. For information about viewing discovered data, see [Viewing Discovered Data](#).

## Related Objects


The Related Objects tab displays the following information about the records associated with the IP address:

- **Name:** The name of the object. For example, if the IP address belongs to a host record, this field displays the hostname. The appliance highlights disabled DHCP objects in gray. A DHCP object can be a DHCP range, fixed address, reservation, host configured for DHCP, or roaming host with an allocated IP address.
- **Type:** The object type, such as DHCP lease, host, A record, and bulk host.
- **Comment:** Information about the object. You can also select the following for display:
- **DNS view:** The DNS view to which the object belongs. You can do the following in this tab:
- Add a resource record. You can select the following from the drop-down list:
  - Host Record—For information, see [Adding Host Records](#).
  - Range—For information, see [Adding IPv4 Address Ranges](#).
  - Fixed Address—For information, see [Adding IPv4 Fixed Addresses](#).
  - Reservation—For information, see [Adding IPv4 Reservations](#).
  - A Record—For information, see [Adding A Records](#).
  - PTR Record—For information, see [Adding PTR Records](#).
- Edit the properties of the selected object. Depending on the type of object, Grid Manager displays the corresponding editor for the object. For example, if the selected object is a fixed address, Grid Manager displays the fixed address editor. When you select a lease object, Grid Manager displays the lease viewer.
- You can also modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#). Delete a selected object or multiple objects.
- When you select a lease object and click the Show Details icon, you can view the lease start and end dates.
- Depending on the object type, you can convert a selected object to one of the following:
  - **Reservation**
  - **Host**
  - **Fixed Address**
- View detailed lease information about the IP address, as described in [Viewing Detailed Lease Information](#).
- Print and export the information in the Related Objects table.

## Audit History

By default, the Audit History tab displays the following information about the last five actions performed on the selected IP:

- **Timestamp:** The day, date, and time of the operation.
- **Action:** The type of operation that was performed by the administrator.
- **Object Type:** The object type of the entry.
- **Object Name:** The name of the object.
- **Admin Name:** The name of the administrator who performed the operation.
- **Message:** The description of the administrative activity.



Note

If you change the IP address of an existing record to a new one in the **IP MAP** tab when the Grid **Audit Logging** is set to *Brief*, then NIOS will not display modification or transition details about the new IP address in this tab. You can only view subsequent modifications and deletions to the new IP address. However, you can view the audit log history and transition details of the old IP address, but you cannot view the initial transition from an old IP address to the new IP address.

## Managing IPv4 Addresses

You can do the following from the IP Map and List panels:

- Add IP addresses to existing hosts. For information, see [Adding IP Addresses to Existing Host Records](#) below.
- Clear unmanaged IP addresses. For information, see [Clearing Unmanaged Data](#) see below.
- Convert objects to other object types. For information, see [Converting Objects Associated with IP Addresses](#).
- Reclaim IP addresses. For information, see [Reclaiming Objects Associated with IPv4 and IPv6 Addresses](#).
- Ping IP addresses. For information, see [Pinging IP Addresses](#).
- Configure and run a network discovery. For information, see [IP Discovery and vDiscovery](#).
- Resolve discovery conflicts. For information, see [Resolving Conflicting Addresses](#).
- Clear discovered data. For information, see [Clearing Discovered Data](#).

### Adding IP Addresses to Existing Host Records

You can add unused and unmanaged addresses, including all their information, to existing host records. When you add an unmanaged address to a host record, the appliance adds the discovered data to the host record. You can select the desired host to which you want to add the unmanaged address.

To add an unmanaged IP address to an existing host record:

1. From the IP Map or List panel, select an IP address, and then click **Add** -> **Add to Existing Host** from the Toolbar.
2. In the *Select Host* dialog box, do the following:
  - In the table, select the host to which you want to add the selected IP address. You can also use the filters or the **Go To** field to narrow down the host list. For information, see [Using Filters](#) and [Using the Go To Function](#).
  - Click the Select icon.  
Grid Manager displays the *Host Record* editor.
3. In the *Host Record* editor, update the host properties.
4. Save the configuration and click **Restart** if it appears at the top of the screen. To close the editor without saving the changes, click the **Close** icon.

### Clearing Unmanaged Data

You can clear the status of unmanaged data at the network and IP address levels. When you clear an unmanaged address, the status of the IP address changes to **Unused**. An unmanaged address is an IP address with a discovered host, is not previously known to the appliance, and does not have an A record, PTR record, fixed address, host address, lease, or is not within a DHCP range. You can change an unmanaged address to a host, a DHCP fixed address, an A record, or a PTR record. You can also clear the unmanaged data associated with the address.

To clear unmanaged data:

1. From the IP Map or List panel, select the IP address for which you want to clear unmanaged data, and then click **Clear** -> **Clear Unmanaged Data** from the Toolbar. You can select multiple IP addresses.
2. In the *Clear Unmanaged data* dialog box, click **Yes**.

### Viewing Router Redundancy Information for Virtual IPs (VIPs)

1. From the **Data Management** tab, click **CSV Job Manager** from the Toolbar.
2. In **CSV Job Manager**, click the **CSV Import** tab and select the import job that you want to delete. Click the Action icon  
≡

and select **Delete** or click the Delete pending job icon.

## Managing IPv6 Networks

In Grid Manager, you can use the IPv6 Net Map (network map) and List panels to manage your IPv6 network infrastructure. After you select a network container from the IPAM tab, Grid Manager displays it in the Net Map panel, by default. The Net Map panel provides a graphical view of your networks and has a number of features that simplify network management. The List panel displays the networks in table format.

You can always switch your view of a network container between the Net Map and List panels. Grid Manager keeps track of which panel you last used. When you select a network container, Grid Manager displays it in the Net Map or List panel, depending on which one you last used. For information about each panel, see [IPv4 Network Map](#) and [IPAM Home](#).

You can use Grid Manager to manage IPv6 networks and their AAAA, PTR and host resource records. You can configure IPv6 networks and track IP address usage in those networks. You can also split and join IPv6 networks, when necessary.

### IPv6 Network Map

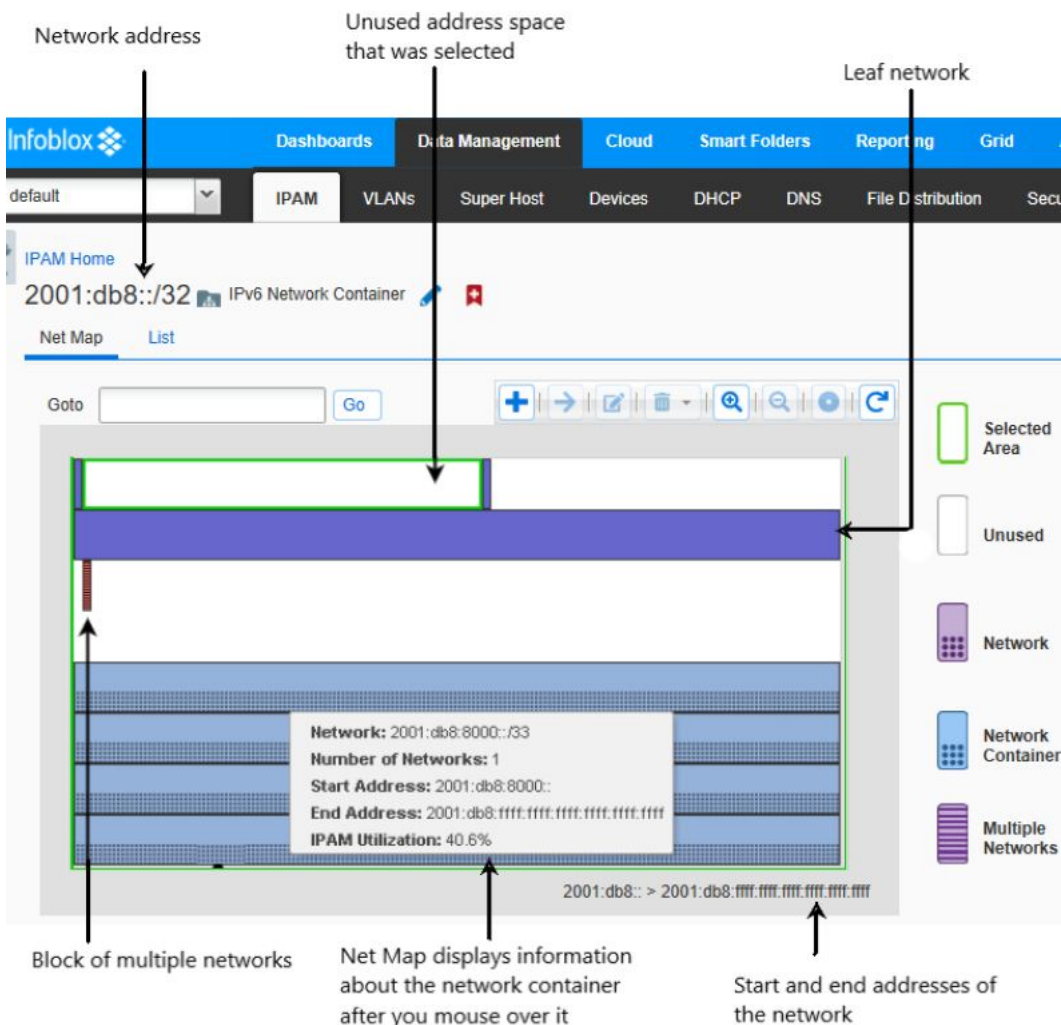
After you select an IPv6 network container from the IPAM tab, Grid Manager displays it in the IPv6 Net Map (network map) panel, by default. Just like the IPv4 Net Map, the IPv6 Net Map provides a high-level view of the network address space. You can use Net Map to design and plan your network infrastructure, and to configure and manage individual networks.

The Net Map panel presents a complete view of the network space, including the different types of networks that are in it and its unused address space. IP addresses that belong to a network are blocked off. Each color-coded block represents a network container, a leaf network, or a block of networks that are too small to be displayed individually in the map. For example, in a /64 or /96 network, networks smaller than /76 or /108 respectively and that are beside each other are represented as a multiple network block. In addition, the fill pattern of the blocks indicates their utilization. Therefore, you can quickly evaluate how many networks are in a network container, their relative sizes, utilization, and how much space you have left.

As you mouse over areas of the map, it displays IP information about the area. Net Map also has a zoom feature that allows you to enlarge or reduce your view of a particular area.

The below figure displays the network map of a 1111::/16 network, which is a network container that has network containers and leaf networks.

*IPv6 Network Map*



### Displaying Network Information

As shown in the figure IPv6 Network Map above, as you mouse over the map, Net Map displays IP information about the area. When you mouse over an unused area, Net Map displays the following information:

- The start and end IP address
- The largest possible network
- The number of /64 networks that can fit in that space

When you mouse over a network container, Net Map displays the following information:

- Network address and netmask
- The first and last IP address of the network
- The number of networks in that block
- IPAM utilization

When you mouse over a network, Net Map displays the following information:

- Network address and netmask
- The first and last IP address of the network

When you mouse over a block of multiple networks, Net Map displays the following information:

- The start and end IP address of that block of networks
- The number of networks in that block



## Zooming In and Out

Use the zoom function to enlarge and reduce your view of a selected area. You can zoom in on any area in your network. You can zoom in on an area until it displays 128 addresses per row, for a total of 1024 addresses for the map. When you reach the last possible zoom level, the Zoom In icon in the Net Map task bar and the menu item are disabled. After you zoom in on an area, you can click the Zoom Controller icon to track where you zoomed in. The Zoom Controller lists all the areas that you zoomed in and updates its list dynamically. You can click an item on the list to view that area again. Click the Zoom Controller again to close it.

To select an area and zoom in:

1. Right-click and select **Zoom In**, or click the Zoom In icon in the Net Map task bar. The pointer changes to the zoom in selector.
2. Select a starting point and drag to the end point. The starting point can be anywhere in the map. It does not have to be at the beginning of a network.  
Net Map displays a magnified view of the selected area after you release the mouse button. As you mouse over the zoomed in area, Net Map displays IP information about it.
3. You can do the following:
  - Select an area and zoom in again.
  - Add a network. If you zoom in on an area and click Add without selecting an open area first, Net Map selects the area where it can create the biggest possible network in that magnified area.
  - Select a network and perform any of the following operations:
    - Edit its properties.
    - Open it to display its IP List.
    - Delete it immediately, or schedule its deletion.
  - Right-click and select **Zoom Out**, or click the Zoom Out icon in the Net Map task bar. Each time you click **Zoom Out**, Net Map zooms out one level and the Zoom Controller is updated accordingly.

## Net Map Tasks

From Net Map, you can create IPv6 networks, and evaluate and manage your network resources according to the needs of your organization. You can do the following:

- Zoom in on specific areas, as described in [Zooming In and Out](#) above.
- Use the **Go to** function to find a network in the current zoom level of Net Map.
- Add a network, as described in [Adding a Network from Net Map](#) below.
- Select a network and view IP address list, as described in [Viewing IPv6 Data](#).
- Select a network and edit its properties, as described in [Modifying IPv4 and IPv6 Network Containers and Networks](#).
- Split a network, as described in [Splitting IPv6 Networks into Subnets](#) below.
- Join networks, as described in [Joining IPv6 Networks](#) below.
- Delete one or multiple networks, as described in [Discovering Networks \(Under Network Insight only\)](#).
- Switch to the List view of the network. For information, see [IPv6 Network List](#) below.
  - When you select one or more networks in Net Map and then switch to the List view, the list displays the page with the first selected network.
  - If you select one or more networks in the List view and then switch to the Net Map view, the first network is also selected in Net Map. Although, if you select a network in the List view that is part of a Multiple Networks block in Net Map, it is not selected when you switch to the Net Map view.

## Adding a Network from Net Map

When you create networks from Net Map, you can view the address space to which you are adding a network, so you can determine how much space is available and which IP addresses are not in use. When you mouse over an open area, Net Map displays useful information, such as the largest possible network that fits in that area. In addition, you can create networks without having to calculate anything. When you add a network, Net Map displays a netmask slider so you can determine the appropriate netmask for the size of the network that you need. As you move the slider, it displays network information, including the total number of addresses. After you select the netmask, you can even move the new network around the open area to select another valid start address.

To add a network from the Net Map panel:

1. Do one of the following:
  - Click the Add icon.  
Net Map displays the netmask slider and outlines the open area that can accommodate the largest network.
  - Select an open area, and then click the Add icon.  
Net Map displays the netmask slider and outlines the largest network that you can create in the open area that you selected.
2. Move the slider to the desired netmask. You can move the slider to the netmask of the largest network that can be created in the open area. You can also move the slider to the smallest network that can be placed in the current zoom level of Net Map.  
As you move the slider, Net Map displays the netmask. The outline in the network map also adjusts as you move the slider. When you mouse over the outline, it displays the start and end address of the network.
3. After you set the slider to the desired netmask, you can drag the new network block around the open area to select a new valid starting address. You cannot move the block to a starting address that is invalid.
4. Click **Launch Wizard** to create the network.  
The *Add Network* wizard displays the selected network address and netmask.
5. You can add comments, automatically create reverse mapping zones, and edit the extensible attributes. (For information, see [Adding IPv6 Networks.](#))
6. Save the configuration and click **Restart** if it appears at the top of the screen. Grid Manager updates Net Map with the newly created network.

## Viewing Network Details

From Net Map, you can focus on a specific network or area and view additional information about it. If you have a network hierarchy of networks within network containers, you can drill down to individual leaf networks and view their IP address usage.

1. Select a network or area.
2. Click the Open icon.
  - If you selected a network container, Grid Manager displays it in the Net Map panel. You can drill down further by selecting a network or open area and clicking the Open icon again.
  - If you selected a block of multiple networks, Grid Manager displays the individual networks in the Net Map panel. You can then select a network or open area for viewing.
  - If you selected a leaf network, Grid Manager displays it in the Network List panel.
  - If you selected an open area, Grid Manager displays an enlarged view of that area in the Net Map panel. This is useful when you are creating small networks in an open area.

## IPv6 Network List

The Network list panel is an alternative view of an IPv6 network hierarchy. For a given network, the panel shows all the networks of a selected network view in table format. A network list displays only the first-level subnets. It does not show further descendant or child subnets. You can open a subnet to view its child subnets. Subnets that contain child subnets are displayed as network containers. If the number of subnets in a network exceeds the maximum page size of the table, the network list displays the subnets on multiple pages. You can use the page navigation buttons at the bottom of the table to navigate through the pages of subnets.

The IPAM home panel displays the following:

- **Network:** The network address.
- **Comment:** Information you entered about the network.
- **IPAM Utilization:** For a network, this is the percentage based on the IP addresses in use divided by the total addresses in the network. You can use this information to verify if there is a sufficient number of available addresses in a network. The IPAM utilization is calculated approximately every 15 minutes.
- **Site:** The site to which the IP address belongs. This is a predefined extensible attribute.
- **Active Users:** The number of active users on the selected network. You can select the following columns for display:
- **Disabled:** Indicates whether the network is disabled.
- **Leaf Network:** Indicates whether or not the network is a leaf network.
- Other available extensible attributes



You can sort the list of subnets in ascending or descending order by columns. For information about customizing tables in Grid Manager, see [Customizing Tables](#).

You can also modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).

---

**Tip:** If you select a network from the list and switch to the Net Map panel, the network is also selected in the network map.

---

### Filtering the Network List

You can filter the network list, so it displays only the networks you need. You can filter the list based on certain parameters, such as network addresses, comments and extensible attributes. When you expand the list of available fields you can use for the filter, note that the extensible attributes are those with a gray background.

### Splitting IPv6 Networks into Subnets

You can create smaller subnets simultaneously within a network by splitting it. You do not have to configure each subnet individually. You can create smaller subnets with larger netmasks. A larger netmask defines a larger number of network addresses and a smaller number of IP addresses.

Note that you cannot split a network that is part of a shared network. To split an IPv6 network:

1. From the **Data Management** tab, select the **IPAM** tab -> *network* checkbox, and then click **Split** from the Toolbar.
2. In the *Split Network* editor, do the following:
  - **Address:** Displays the network address. You cannot modify this field.
  - **Net mask:** Specify the appropriate netmask for each subnet.
  - **IPv6 Prefix Collector Network:** If you split a network with prefix delegations that are not tied to specific addresses, specify the network in which all prefix delegations are assigned. If you leave this field blank, the server assigns all prefix delegations that are not tied to specific addresses to the first network.
  - **Immediately create:** Select one of the following:
    - **Only networks with ranges and fixed addresses:** Adds only the networks that have DHCP ranges and fixed addresses.
    - **All possible networks:** Adds all networks that are within the selected netmasks. You can select this option only when you increase the CIDR by 8 bits.
  - **Automatically create reverse-mapping zone:** Select this checkbox to have the appliance automatically create reverse-mapping zones for the subnets. This function is enabled if the netmask of the network is a multiple of four, such as 4, 12 or 16.
3. Click **OK**.

### Joining IPv6 Networks

Joining multiple networks into a larger network is the opposite of splitting a network. You can select a network and expand it into a larger network with a smaller netmask. A smaller netmask defines fewer networks while accommodating a larger number of IP addresses. Joining or expanding a network allows you to consolidate all of the adjacent networks into the expanded network. Adjacent networks are all networks that fall under the netmask of the newly-expanded network.

To join or expand a network:

1. From the **Data Management** tab, select the **IPAM** tab -> *network* checkbox, and then click **Join** from the Toolbar.
2. In the *Join Network* editor, do the following:
  - **Address:** Displays the network address. You cannot modify this field.
  - **Netmask:** Enter the netmask of the expanded network.

- **Automatically create reverse-mapping zone:** Select this checkbox to configure the expanded network to support reverse-mapping zones. The appliance automatically creates reverse-mapping zones only if the netmask is between /4 through /128, in increments of 4 (that is, /4, /8, /12, and so on until /128).
3. Click **OK**.

## Viewing IPv6 Data

To the configured IP addresses in an IPv6 network:  
by selecting an IPv6 leaf network from the Network List panel

- For a leaf network that is not in a network container, from the **Data Management** tab, select the **IPAM** tab, and then click the IPv6 network you want to view.
- For a leaf network that is in a network container, from the **Data Management** tab, select the **IPAM** tab → *network\_container* -> *network*.

Grid Manager lists the configured IPv6 addresses. You can export and print the list. It displays the following information about each IP address:

- **IP Address:** The name of the IPv6 DHCP object, which can be a DHCP range, fixed address, host configured for DHCP, or a roaming host with an allocated IP address.
- **Name:** The name of the record associated with the IP address.
- **DUID:** The DHCP Unique Identifier (DUID) of the device that was assigned the IP address.
- **Status:** The status of the IPv6 object, such as Used or Unused.
- **Type:** The object type associated with the IP address, such as **AAAA record**, **IPv6 Fixed Address**, or **Unmanaged**.
- **Usage:** Indicates whether the IPv6 address is configured for DNS or DHCP.
- **Exclude:** (*Applies only with Network Insight*) Denotes whether the IP is excluded from discovery.
- **Lease State:** (*Applies only with Network Insight*) The lease state of the record, such as Active.
- **User Name:** The name of the user who received the lease for the IP address.
- **First Discovered:** (*Applies only with Network Insight*) The date and timestamp of the first occasion that NIOS discovered the IP address.
- **Last Discovered:** (*Applies only with Network Insight*) The date and timestamp of the last occasion that NIOS discovered the IP address.
- **OS:** The operating system of the IP.
- **NetBIOS Name:** The returned NetBIOS name from the last discovery.
- **Device Type(s):** Shows the device type for the device associated with the IP address.
- **Fingerprint:** The name of the DHCP fingerprint or vendor ID of the network device that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#).
- **Comment:** Displays comments about the record.
- **Site:** The site to which the IP address belongs. This is a predefined extensible attribute.

You can display all available extensible attributes. You can also sort the list of IP addresses in ascending or descending order by **IP Address** only.

You can drill down further and view the records associated with an IP address. To view the associated records of an IP address, select it and Grid Manager displays information about the IP address in the **Related Objects** and **Audit History** tabs.

### Related Objects

Grid Manager displays the following information about the records associated with the IP address:

- **Name:** The record name. For example, if the IP address belongs to a host record, this field displays the hostname.
- **Type:** The object type. For example, AAAA Record, PTR Record, Host Record, IPv6 Fixed Address.
- **Comment:** Additional information that was entered in the record about the IP address.

### Audit History

Grid Manager displays the following information about the last five actions performed on the selected IP:

- **Timestamp:** The day, date, and time of the operation.
- **Action:** The type of operation that was performed by the administrator.
- **Object Type:** The object type of the entry.
- **Admin Name:** The name of the administrator that performed the operation.
- **Message:** Description of the administrative activity.

### Filtering the IP Address List

You can filter the IP address list, so it displays only the IP addressees you need. You can filter the list based on any combination of extensible attributes and the parameters displayed in the IP address list, such as usage and type. When you expand the list of available fields you can add to the filter, note that the extensible attributes are those with the gray background.

## Managing IPv4 and IPv6 Addresses

Grid Manager uses IP addresses as the entry point to the data set containing Infoblox host, DNS, DHCP, and other information related to that address. You can view the data, modify it, assign extensible attributes to the objects associated with the address, and convert DHCP lease types, such as changing a currently active dynamic lease to a fixed address or host record.

You can view and manage IPv4 address data in the IP Map panel, and view and manage IPv4 and IPv6 data in the IP List panel. You can do the following for IPv4 and IPv6 data from the IP List panel:

- Convert objects to other object types.
- Reclaim IP addresses.
- Ping IP addresses.
- Clear DHCP leases.

You can also print and export in CSV format the information displayed in any panel that supports these functions.

### Converting Objects Associated with IP Addresses

The NIOS appliance provides a simple mechanism for converting unmanaged IP addresses to resource records, such as host records and A or AAAA records. You can also convert the active lease of a dynamically assigned IPv4 or IPv6 address to a fixed address or host, and convert an IPv4 lease to an IPv4 reservation. Using the conversion mechanism, you can keep the existing information of a network device during the conversion.

The appliance supports the following conversions for IPv4 objects:

- DHCP leases to fixed addresses, reservations, or host records
- Fixed addresses to reservations or host records
- Unmanaged addresses to host records, A records, PTR records, or fixed addresses
- A records to host records
- PTR records to host records

The appliance supports the following conversions for IPv6 objects:

- DHCP leases to fixed addresses or host records
- Fixed addresses to host records
- AAAA records to host records
- IPv6 PTR records to host records



#### Note

You cannot convert unmanaged IP addresses or leases served by Microsoft DHCP servers to host records.

## Converting DHCP Leases

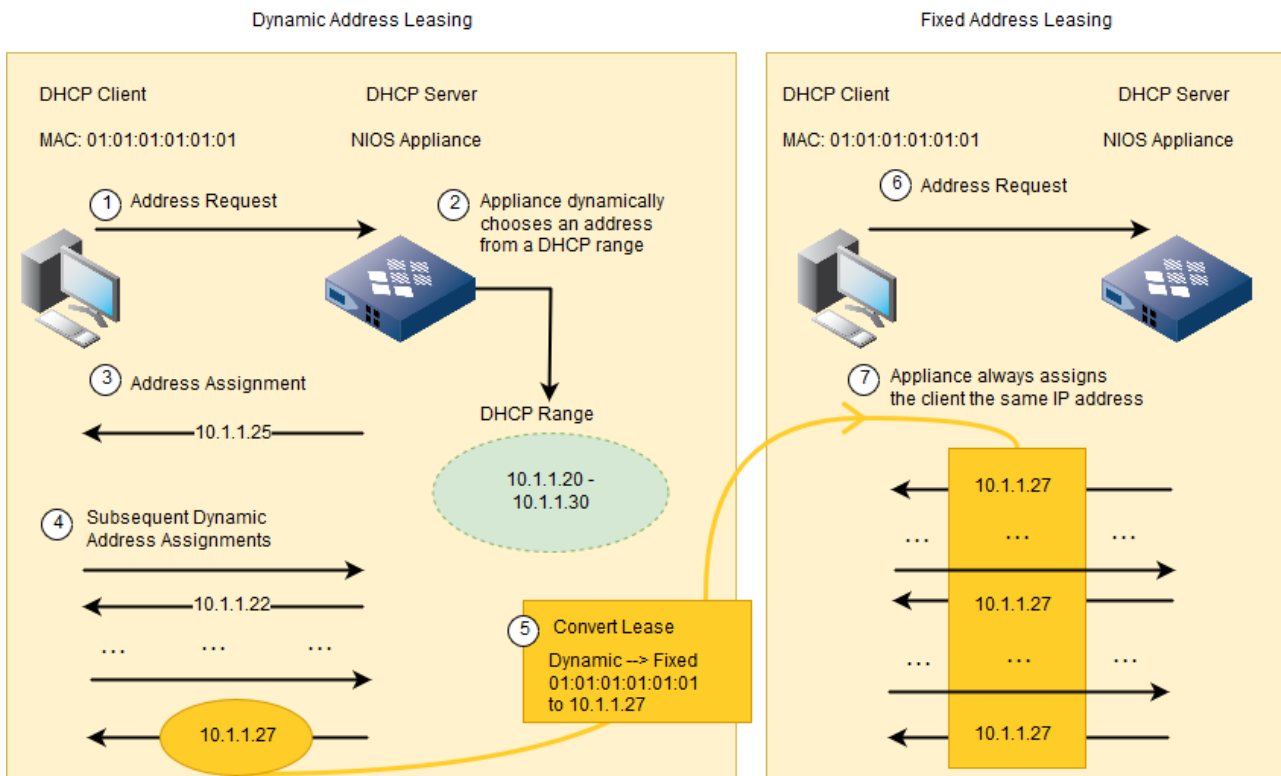
To create a fixed address, you bind an IPv4 address to a MAC address or an IPv6 address to a DUID. You can make that binding by converting an active dynamically leased address to a fixed address. The lease conversion transforms the temporary binding between the IPv4 address and MAC address or the IPv6 address and DUID in the dynamic lease to a persistent one. The lease must be active so that the NIOS appliance has an IPv4-to-MAC address or IPv6-to-DUID binding to convert into a fixed address.

The appliance uses the following rules when converting a DHCP lease:

- If an IPv4 DHCP lease is converted to a fixed address, the appliance copies the client identifier to the fixed address, based on information in the lease. If the appliance finds the client identifier in the lease information, the appliance includes it when it creates the host. If it finds the MAC address, the appliance includes it when it creates the host. If it finds both, the appliance includes only the MAC address (default) when it creates the host.
- If an IPv6 DHCP lease is converted to a fixed address, the appliance copies the DUID to the fixed address.
- If you try to convert an IPv4 DHCP lease or a fixed address with a client identifier, not a MAC address, to a host, the appliance displays an error message in the host editor. This ensures that you do not attempt this operation and lose the data.
- You cannot create two IPv4 fixed addresses with the same client identifier or MAC address in the same network. You cannot create two IPv6 fixed addresses with the same DUID in the same network.
- If the appliance receives a second IPv4 DHCP request with the same client identifier, it provides the same fixed IP address if the lease is still binding.

The below figure illustrates converting a dynamic IPv4 lease to a fixed lease.

*Converting a Dynamic IPv4 Lease to a Fixed Lease*



An advantage of converting an active dynamic lease is that you do not need to learn the MAC address or DUID of the device to which you want to assign an IP address and manually enter it in the fixed address configuration.

An IPv4 reservation is an address that you exclude from DHCP use because you intend to configure that address manually on a device, such as a firewall, router, or printer. You can also convert an IPv4 fixed address or a dynamic address with an active lease to a reservation.

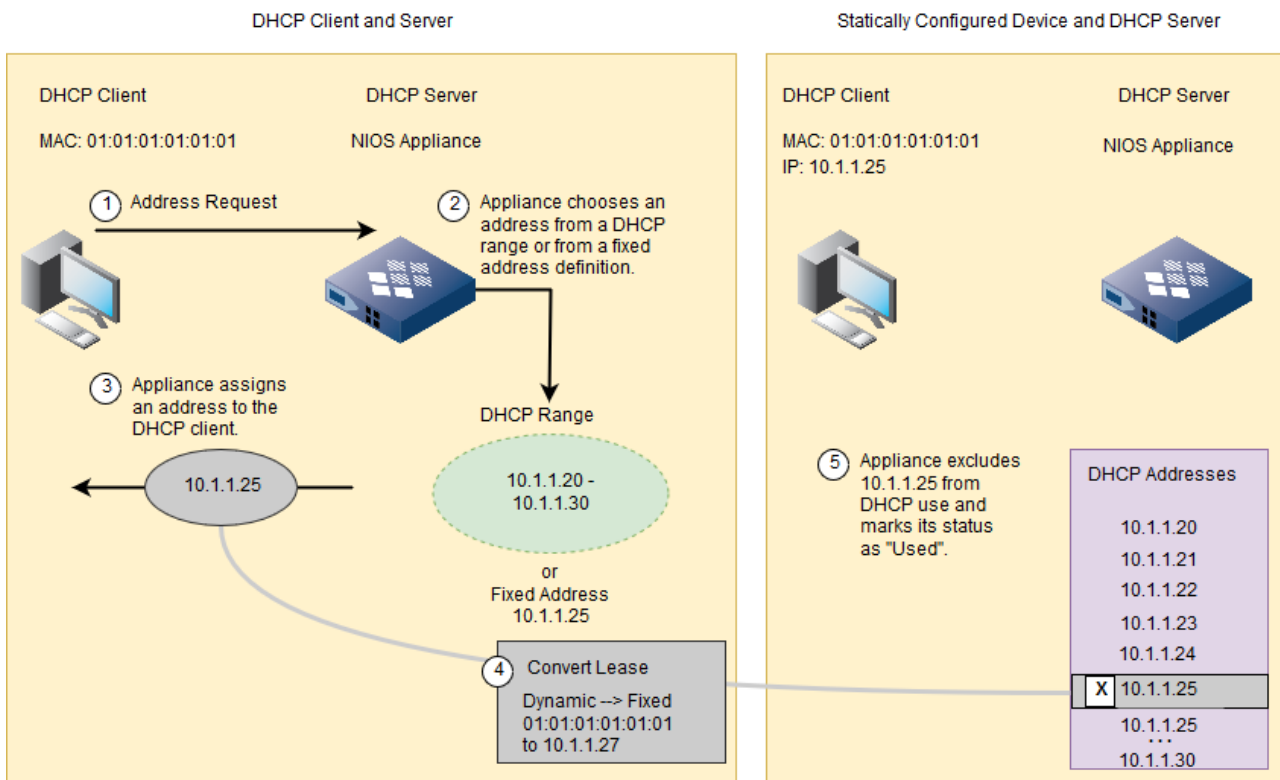
When you convert an address in a DHCP range to a reservation, you reduce the total number of dynamically assignable

addresses in that range by one. Correspondingly, this reduces the number of allocated addresses needed to exceed a high or low watermark threshold for that range.

**Note**  
To return an IP address to its place in a DHCP range after converting it from an active dynamic lease to a fixed address, reservation, or Infoblox host, delete the fixed address, reservation, or host to which you previously converted the IP address. The IP address then becomes part of the DHCP range to which it first belonged.

You can convert IPv4 fixed addresses to reservations, as shown in the below figure.

*Converting an IPv4 Dynamic Lease or Fixed Address to a Reservation*  
DHCP Client and Server Statically Configured Device and DHCP Server



To convert an object:

1. From the IP Map, select an IPv4 address or from the IP List panel, select an IPv4 or IPv6 address.
2. In the Related Objects tab, select the checkbox of the object, and then click **Convert** from the Toolbar or navigation bar.
3. Select the object type to which you want to convert the object. Grid Manager displays the corresponding editor for the object type.
4. For all IPv4 conversions, Grid Manager populates the discovered information in the corresponding editor. Depending on the type of conversion, do one of the following:
  - For host record conversions, see [About Host Records](#).
  - For IPv4 reservation conversions, see [Modifying IPv4 Reservations](#).
  - For fixed address conversions, see [Modifying IPv4 Fixed Addresses](#).
  - For A record conversions, see [Modifying A Records](#).
  - For PTR record conversions, see [Modifying PTR Records](#).

**Note**

When you select an object for conversion, Grid Manager displays only the available conversion types for the object. You must save the changes in the editor for the conversion to take place.

## Reclaiming Objects Associated with IPv4 and IPv6 Addresses

You can use the reclaim IP function to delete all objects, except the active DHCP lease, that are associated with a selected IP address. To delete a DHCP lease, use the clear lease function as described in [Clearing Active DHCP Leases](#) below. When you reclaim an IP address, Grid Manager deletes the associated objects and puts them in the Recycle Bin, if enabled. You can reclaim any used and unmanaged IP addresses. You can also select multiple IP addresses for this function. After you reclaim an IP address, the address status changes to Unused. You can then reassign the IP address to other objects. For example, when you reclaim a fixed address, Grid Manager deletes the fixed address object and puts it in the Recycle Bin. When you reclaim an IP address that is associated with a host record and the address is the only address in the host, Grid Manager deletes the host record.

Grid Manager deletes all the objects that are associated with the selected IP addresses and puts them in the recycle bin, with the following exceptions:

- When you reclaim IP addresses that are in a DHCP range, all the objects that are associated with the IP addresses are deleted and the IP addresses remain in the DHCP range.
- When you select an IP address that is part of a host record, only the selected IP address is deleted from the host. However, if the selected address is the only address in the host, Grid Manager deletes the host record.

Grid Manager does not reclaim the following:

- Unused IP addresses
- Bulk hosts

To reclaim an IP address:

1. From the IP Map or List panel, select the IP address you want to reclaim, and then click **Reclaim** from the Toolbar. You can select multiple IP addresses.
2. In the *Delete Confirmation* dialog box, click **Yes**.

Grid Manager puts the deleted objects in the Recycle Bin, if enabled.

## Pinging IP Addresses

You can find out whether an IP address is accessible and active by pinging the address. Grid Manager sends a packet to the selected IP address and waits for a reply when you ping the address. You can ping individual IP addresses from the IP Map and IP List panels. You can ping all IP addresses from the IP Map panel and all IP addresses on the selected page from the IP List panel.

To ping an IPv4 or IPv6 address:

- From the IP Map or IP List panel, select the IP address that you want to ping, and then click **Ping** from the Toolbar.

To ping all IPv4 addresses:

- From the IP Map panel, click **Multi-ping** from the Toolbar. Grid Manager pings all IP addresses displayed in the IP Map panel and displays the ping status in the panel. When the ping or multi-ping is complete, the status bar displays the number of active IP addresses detected through the ping. To close the ping status bar, click the Close icon.
- From the IP List panel, click **Multi-ping** from the Toolbar. Grid Manager pings all IP addresses visible on the selected page. When the ping or multi-ping is complete, the status bar displays the number of active IP addresses detected on the selected page. To close the ping status bar, click the Close icon.

## Clearing Active DHCP Leases

A DHCP lease specifies the amount of time that the DHCP server grants to a network device the permission to use a particular IP address. You may sometimes need to terminate an active lease. The following are some of the reasons for clearing active DHCP leases:

- When a network device is moved to another network.
- Reset a DHCP lease to fix other problems.

In Grid Manager, you can select multiple IP addresses and clear their active DHCP leases. To clear an active lease:

1. From the IP Map or List panel, select the IP address for which you want to clear a DHCP lease, and then click **Clear** -> **Clear Lease** from the Toolbar. You can select multiple IP addresses.
2. In the *Clear DHCP Lease Confirmation* dialog box, click **Yes**.

## Configuring Thresholds for IPAM Utilization

You can define thresholds for IPAM utilization in a network or network container and configure the appliance to send SNMP trap and email notifications to a designated destination when IPAM utilization in a network or network container crosses the configured threshold. IPAM utilization for a network is the percentage based on the IP addresses in use divided by the total addresses in the network and for a network container that contains subnets, it is the percentage of the total address space defined within the container regardless of whether any of the IP addresses in the subnets are in use. The appliance sends an SNMP trap and email notification only once when the IPAM utilization in a network or network container exceeds the Trigger value and when it drops below the Reset value. The default Trigger value is 95% and the default Reset value is 85%. The IPAM utilization notifications are sent for IPv4 networks and IPv4 network containers only. The appliance updates the IPAM utilization data immediately for a network container, but for a network it is updated every 15 minutes.

You can define thresholds for IPAM utilization at the Grid level and network level. The appliance applies the settings hierarchically in a parent-child structure. By defining thresholds at a higher level, all networks can then inherit the same settings and you do not have to redefine them for each network. For example, if you set the thresholds for IPAM utilization at the Grid level, then the settings applies to all the network containers and networks in any network view. However, if you override these settings at the network container or network level, then the settings applies to that network or network container and any network within that network or network container in the same network view. If you set the thresholds for an individual network, then it overrides settings at a higher level.



### Note

Infoblox recommends that you do not enable SNMP traps and email notifications for IPAM utilization during an upgrade, because if you have configured notifications you may have to unconfigure them during an upgrade.

You can configure the thresholds for IPAM utilization at the Grid level and override them at the network level. To configure the IPAM utilization thresholds at the Grid level, see [Defining Threshold for Traps](#).

To configure the IPAM utilization thresholds for a IPv4 network, network container, or network template, complete the following:

1. **Network Container:** From the **Data Management** tab, select the **IPAM** tab -> *network\_container* checkbox, and click the Edit icon.  
**Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* checkbox, and click the Edit icon.  
**Network Template:** From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *DHCP\_template* checkbox, and click the Edit icon.
2. In the editor, click **Toggle Advanced Mode**, select the **IPv4 IPAM Utilization Notification** tab, and then complete the following:
  - **IPAM Utilization Notification:** Click **Override** to override the inherited property, and specify the following:
    - **Enable SNMP Notifications:** Select this for the appliance to send an SNMP trap to the trap receiver that you define for the Grid when IPAM utilization crosses the configured threshold.
    - **Enable Email Notifications:** Select this for the appliance to send an email notification to a designated destination if IPAM utilization crosses the configured threshold.
  - **IPAM Threshold Settings:** Click **Override** to override the inherited property, and specify the following:
    - **Trigger:** Enter a Trigger value between 0 to 100. The appliance sends an SNMP trap and—if configured to do so—sends an email notification to a designated destination when the IPAM utilization exceeds the Trigger value. The default Trigger value is 95.



- **Reset:** Enter a Reset value between 0 to 100. The appliance sends an SNMP trap and —if configured to do so—an email notification to a designated destination when the IPAM utilization drops below the Reset value. The default Reset value is 85.
  - **Email Addresses:** Click **Override** to override the inherited property. Click the Add icon, and then enter an email address to which you want the appliance to send email notifications when IPAM utilization for the network or network container crosses the configured threshold. You can create a list of email addresses.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Viewing Identity Mapping Information

You can view user information associated with networks, end-host devices, Active Directory domains, routers and switches when you enable the **Identity Mapping** feature on the appliance. Access to user information related to networks and devices help network administrators to understand how the network resources are consumed and by whom. Each network user being mapped can use different devices to access their network environment. So using the identity mapping feature and synchronizing servers, such as Microsoft servers and Cisco ISEs, on the Infoblox appliance provide visibility of user interaction with their environments. By enabling this feature, you can monitor domain users, the IP addresses they log on to, the login status, and the time duration of their current status in the **IPAM** tab. For information about Identity Mapping for Active Directory users, see [Configuring Identity Mapping](#) and how to collect about user and device information from Cisco ISEs, see [Configuring Cisco ISE on NIOS](#). You can generate the user login history report to monitor user login activities in a given time frame. For information, see [User Login History Report](#). You can do the following in the **Network Users** tab:


- View active network users, as described in [Viewing Active Network Users](#) below.
- View user login history, as described in [Viewing Network Users Login History](#) below.

## Enabling Identity Mapping

You must first enable the Identity Mapping feature to view user information of a network or device. Complete the following to enable this feature:

1. From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties Editor*, select the **General** tab -> **Advanced** tab, and then complete the following:
  - **Enable Network Users feature:** Select this checkbox to enable the Identity Mapping feature on the appliance. Note that the **Network Users** tab is available only after you enable this feature. Note on an Infoblox appliance, the **Enable Network Users Feature** option is disabled by default for all new installations.
3. Save the configuration.

## Viewing Active Network Users

You can view all the users who are currently active on a network in the **Active Users** tab. Using the Action icon , you can do the following in the **Active Users** tab:

- **Go To IPAM IP Address Details:** Select this to open the IPAM Home page to view the network address of the device. This option is greyed out for devices that have an IP address and is not part of an IPAM network.
- **Go To DHCP Network Details:** Select this to open the **DHCP > Networks** tab to view the network address of the device. This option is greyed out for devices that have an IP address and not managed by the Grid.
- **Go To IPAM Network Details:** Select this to open the IPAM Map for the selected user. The page shows network information in graphical format. This option is greyed out for devices that have an IP and not managed by the Grid.

To view active users:

1. From the **Data Management** tab, select the **Network Users** tab > **Active Users** tab.

or

You can use the Action icon

in the following tabs to open the *Active Users* dialog box:

**IPAM** tab: From the **Data Management** tab, select the **IPAM** tab, click the Action icon



next to the respective network and select **Show Active Users**.

**DHCP** tab: From the **Data Management** tab, select the **DHCP** tab > **Networks** tab, click the Action icon

next to the respective network and select **Show Active Users**.

**Cloud** tab:

- In the **Networks** tab, click the Action icon

next to the respective network and select **Show Active Users**.

- In the **VMs** tab, click the Action icon

and select **Show Active Users**.

**DNS** tab: From the **Data Management** tab, select the **DNS** tab > **Zones** tab > **Records** tab, click the Action icon

and select **Show Active Users**.

The **Active Users** tab or *Active Users* dialog box displays the following information:

- **User Name**: Displays the logon name of the user. When the same user logs in to the domain from multiple clients, entry for each IP address is displayed separately. If multiple users logs in to the same domain, entry for each user is listed separately.
- **Domain**: The name of the domain.
- **First Seen**: The timestamp when the user logged in to the network for the first time.
- **IP Address**: The IP address of the client.
- **Data Source**: The IP address of the Microsoft server or the API method.
- **Data Source IP Address**: Displays the source from which the data is collected. It can be Cisco ISE, Microsoft server or the API method.
- **Last Seen**: The timestamp when the user was last seen accessing the network.
- **Last updated**: Displays the timestamp when the user information was last updated.

## Viewing Network Users Login History

You can view the login history of end-host devices, networks, and Active Directory domain users. You must first enable the identity mapping feature to view user login information. For information about enabling Identity Mapping feature, see [Enabling Identity Mapping](#) above.

To view network user login history:

1. From the **Data Management** tab, select the **Network Users** tab -> **User History** tab. Grid Manager displays the following information:
  - **User Name**: The logon name of the user. When the same user logs in to the Active Directory domain from multiple clients, entry for each IP address is displayed separately. If multiple users logs in to the same Active Directory domain, entry for each user is listed separately.
  - **Domain**: Name of the Active Directory domain.
  - **First Seen**: The timestamp when the user logged in to the Active Directory domain for the first time.
  - **Log Out Time**: Displays the log out time of the user.
  - **IP Address**: The IP address of the client.
  - **Data Source**: The IP address of the Microsoft server or the API method.
  - **Status**: Displays the status of the user. The status can be one of the following: **Active** (logged in), **Logged Out**, and **Timed Out**.
    - **Active**: The user is logged in and active.
    - **Logged Out**: The user has logged out of the system.
    - **Timed Out**: The user is logged in but has been idled for a certain period of time. The default is two hours. You can configure this time interval, as described in [Configuring Active User Timeout Session](#) below.
  - **Last Seen**: The timestamp when the user was last seen accessing the network.
  - **Last updated**: The timestamp when the user information was last updated.

## Configuring Active User Timeout Session

You can configure the amount of time that an active session of a user changes to timed out. When the idle session time is reached, the user status changes to inactive status. The default idle time is 2 hours. You can change it to minutes, hours, or days. The user status can be one of the following: **Active**, **Logged Out**, and **Timed Out**.

To configure active user timeout interval:

1. From the **Grid** tab -> **Grid Manager** tab, expand the Toolbar and click **Grid Properties** -> **Edit**.
2. Select **Microsoft Integration** tab in the *Grid Properties Editor* and complete the following in the **Basic** tab:
  - **Assumed Network Users Time Out**: Specify the time period after which the user status changes to **Timed Out**. Select the time period in minutes, hours, or days from the drop-down list.
3. Save the configuration.

## IP Discovery and vDiscovery

This section provides information about the Infoblox IP discovery and vDiscovery, and how you can use them to detect, collect, and manage information about active hosts in predefined networks as well as virtual entities in private, public, and hybrid clouds managed through CMPs (Cloud Management Platforms) such as VMware vCenter servers and vSphere Hypervisor, OpenStack, and AWS (Amazon Web Services). It also explains how to integrate and view discovered data from the NetMRI appliances.

This chapter includes the following topics:

- [About Discovery](#)
- [IP Discovery Process](#)
- [Supported IP Discovery Methods](#)
- [Guidelines Before Starting a Discovery](#)
- [Configuring IP Discovery](#)
- [Configuring vDiscovery Jobs](#)
- [Viewing Discovered Data](#)
- [Managing Discovered Data](#)
- [Integrating Discovered Data from NetMRI](#)

## About Discovery



### Note

The discovery features described in this chapter apply to NIOS Grid deployments that do not use the Discovery license and its accompanying features under Network Insight. Network Insight provides the ability to discover, query, and catalog routed and switched networks and the devices within them, including infrastructure routers, enterprise switches, security devices such as firewalls, wireless access points, end host computer systems, and more. For more information about Network Insight, see [Infoblox Network Insight](#).

Infoblox provides IP discovery for detecting and obtaining information about active hosts in predefined networks, and vDiscovery for discovering virtual entities and interfaces (such as vSwitch and vRouter) in private, public, and hybrid clouds managed through CMPs (Cloud Management Platforms) such as VMware vCenter servers and vSphere Hypervisor, OpenStack, and AWS (Amazon Web Services). You can configure multiple discovery tasks on one or more discovering members.

- IP discovery: You execute an IP discovery (**Data Management** tab -> **IPAM** tab -> **Discovery** from the Toolbar) to detect active hosts on specified networks in a network view. You can configure the appliance to perform an IP discovery using one of the following protocols: ICMP (Internet Control Message Protocol), NetBIOS (Network Basic Input/Output System), and TCP (Transmission Control Protocol). For more information, see [Supported IP Discovery Methods](#). You can start an IP discovery immediately after you configure it, schedule it for a later date

and time, or configure a recurring discovery based on a recurrence pattern. For information about how to configure an IP discovery, see [IP Discovery and vDiscovery](#).

- vDiscovery: This is an extension to the former VM discovery, in which the NIOS appliance only discovers virtual entities on VMware vCenter and vSphere servers. A vDiscovery job (from the **Data Management** tab -> **IPAM** tab -> **vDiscovery** from the Toolbar, or **Cloud** tab -> any sub tab -> **vDiscovery** from the Toolbar) now detects virtual entities and interfaces in private, public, and hybrid clouds that are managed through VMware vCenter servers and vSphere Hypervisor, OpenStack, Azure, or AWS. You can define vDiscovery jobs through the *vDiscovery Job* wizard and manage all configured vDiscovery jobs through the **vDiscovery Job Manager**. Note that for a specific vDiscovery job, NIOS synchronizes successive discovered data (not the associated NIOS objects) with the data in the targeted CMP. For example, if you change the IP address of a VM, this information is reflected in the next discovery of the same vDiscovery job. If you terminate a VM, the VM is deleted from the NIOS database. If you delete certain information on the CMP, the respective discovered data is removed from the NIOS database. Be aware that if you change the parameters of a vDiscovery Job, the last discovered data from this job will be automatically cleaned up so that the appliance can continue to synchronize data from one discovery to the next. If you do not want to lose discovered data for a specific vDiscovery job, you should create a new vDiscovery job for this new collection instead of modifying the current job. For information about how to configure vDiscovery jobs for specific CMPs and how to manage them, see [Configuring vDiscovery Jobs](#) and [Managing vDiscovery Jobs](#).



#### Note

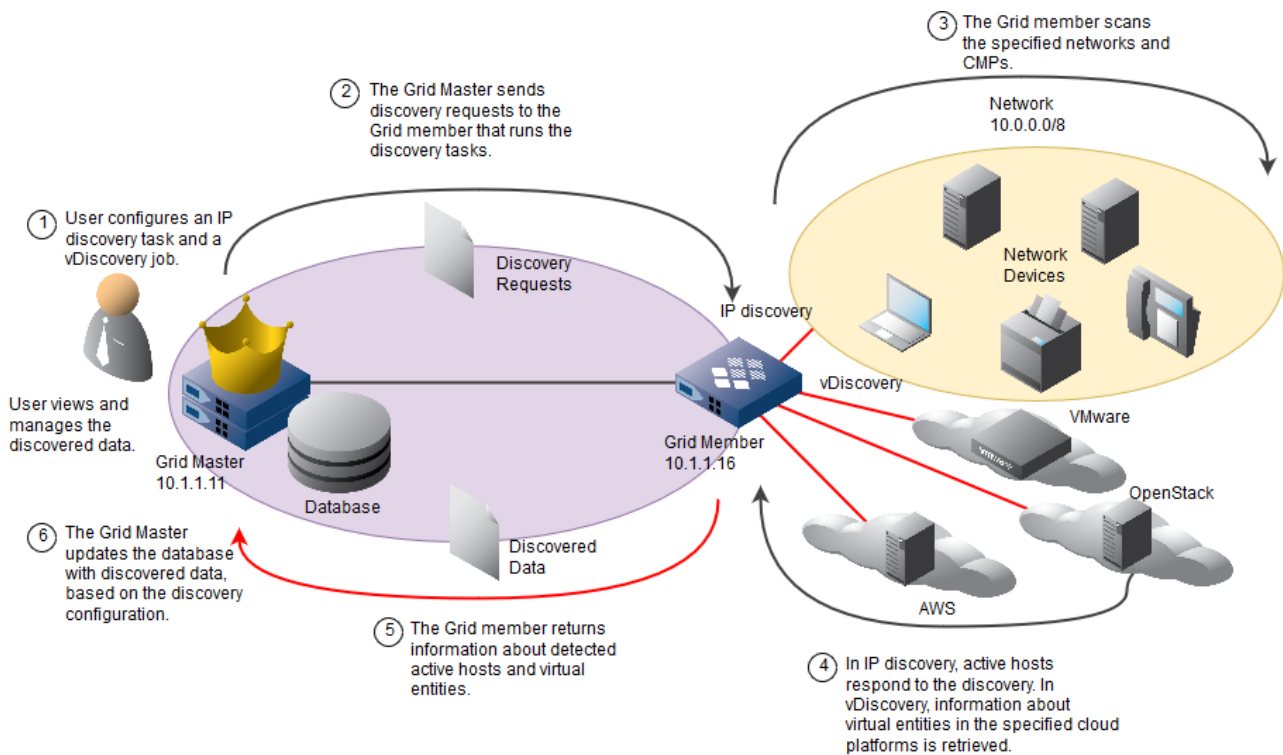
For new installations, an IP discovery task is automatically created by default. You can choose to disable the IP discovery after you have set up your appliance. However, you must configure and manually schedule vDiscovery jobs in order for the appliance to detect and collect information about virtual entities in the clouds. When you upgrade from a previous NIOS release to NIOS 7.2 and later, former VM Discovery tasks are divided into separate vDiscovery jobs based on the server endpoints defined in the VM Discovery tasks. All new vDiscovery jobs inherit the same discovery schedule from the old tasks. You must manually enable the new vDiscovery schedules in order for the appliance to perform vDiscovery jobs. For information about how to enable the vDiscovery schedule, see [Scheduling vDiscovery Jobs](#).

After a discovery, the appliance updates the database with the discovered data based on the discovery configuration. For example, you can configure the appliance to merge newly discovered data, consolidate managed data, or update unmanaged data. The appliance also identifies unmanaged and conflict data after a discovery. Unmanaged data is discovered data that is not configured for DNS or DHCP and has no associated NIOS objects. Conflict data is discovered data that is configured for DNS or DHCP and has associated NIOS object or objects, but certain key values are different than those in the NIOS database. For information about guidelines the appliance uses to update discovered data, see [Guidelines Before Starting a Discovery](#) and [Guidelines for Configuring vDiscovery Jobs](#).

Grid Manager displays discovered data in the **Discovered Data** section of the IP address properties panel when you drill down to individual IPs. For information about how to view and manage discovered data, see [Viewing Discovered Data](#) and [Managing Discovered Data](#). The appliance records admin operations in the audit log and discovery operations in the syslog.

The figure High-Level Discovery Process below shows a high-level perspective of the discovery processes. You can configure and initiate an IP discovery from the *Discovery Manager* wizard and a vDiscovery from the *vDiscovery Job* wizard. You must first select a Grid member that runs the discovery tasks. After you configure an IP discovery task and a vDiscovery job, the Grid Master sends the discovery requests to the selected member. Based on the configuration of the discovery tasks, the selected member runs the discovery and collects information about discovered hosts and virtual entities from the specified networks and cloud platforms. The Grid member then reports the discovered results to the Grid Master. Based on the discovery configuration, the Grid Master updates the database with discovered data.

*High-Level Discovery Process*



### Administrative Permissions

You can initiate a discovery and manage discovered data based on your administrative permissions. You must have read/write permission to "Network Discovery" to initiate and manage IP discovery and vDiscovery. You must have at least read-only permission to "All Tenants" and "All Network Views" to view discovered data in the **VMs (by IP Address)** tab in the **Cloud** tab. To take actions on discovered data, such as resolving conflicts or clear unmanaged data, you must have read/write permissions. For information about how to configure admin permissions, see [About Administrative Permissions](#).

Following are permission guidelines for initiating and controlling a discovery:

- Superusers can initiate and control a discovery on all networks and CMPs.
- Administrators with read/write permission to "Network Discovery" can initiate and control a vDiscovery job or an IP discovery. For IP discovery, only the objects with IP addresses to which the administrators have read/write permission are updated to include the discovered data.

After a discovery is completed, the following permission guidelines apply to viewing and managing discovered data:

- Superusers can view and manage all discovered data.
- Administrators with read/write permission to networks can view all discovered data. They can also add unmanaged data to existing hosts, and resolve IP address conflicts.
- Only administrators with read/write permission to a DNS zone or specific record type can convert unmanaged data to a host, fixed address, reservation, A record, or PTR record.
- Administrators with read-only permission to networks can only view discovered data. They cannot change any discovered data.

### IP Discovery Process

Once an IP discovery starts, the Grid member reports the discovery status, such as **Completed**, **Running**, **Paused**, **Stopped**, or **Error**, in the *Discovery Manager* wizard and the *Discovery Status* widget on the Dashboard. In the *Discovery Status* widget, Grid Manager reports the time when the discovery status was last updated and the numbers of each type of discovered data. For information, see [Managing Discovery](#).

When an IP discovery starts, the appliance divides the IP addresses in a network into chunks, with each chunk containing 64 contiguous IP addresses. The discovery process probes each IP address in parallel and in ascending

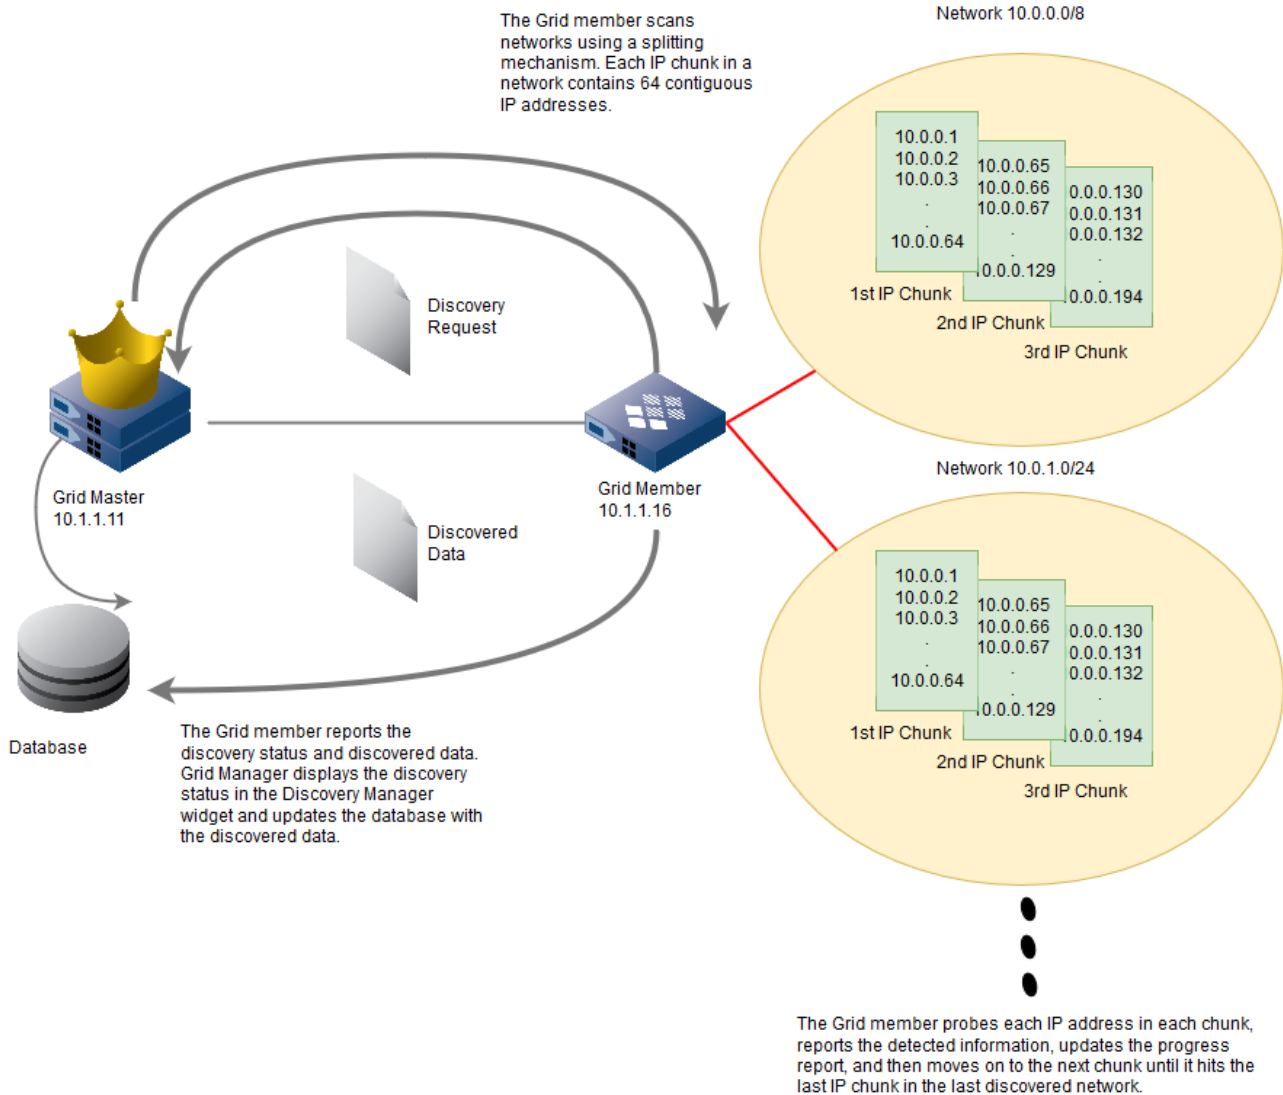
order, reports the detected information, updates the progress report, and then moves on to the next chunk until it hits the last chunk of IP addresses. The appliance then updates the database with the discovered data.

An IP discovery scans the selected networks in the order the networks appear in the *Discover Manager* wizard.

You can configure discovery processes on the same network, but the same configuration cannot be shared between two discovery processes.

The below figure illustrates how an IP discovery works.

### IP Discovery Process



### Supported IP Discovery Methods

When you perform an IP discovery, you can choose one of the following discovery methods:

- ICMP
- NetBIOS
- TCP
- Full

These methods actively scan predefined networks and probe IP addresses. The Grid member listens for responses from the IP addresses as proof of activities. The IP discovery scans through the specified network ranges and probes IP addresses (except for the network, broadcast, and multicast address types) in each network, including the /31 and /32

subnets. Note that the possible addresses in the /31 and /32 subnets can be used only as source addresses for point-to-point links. In these cases, there are no broadcast or network addresses in the /31 and /32 subnets, and the appliance can discover source addresses in these subnets.

## ICMP

This method detects active hosts on a network by sending ICMP echo request packets (also referred to as pings) and listening for ICMP echo responses. The ICMP discovery is a simple and fast discovery that detects whether an IP address exists or not. It returns only the IP address and MAC address (only if the Grid member running the discovery is on the same discovered network) of a detected host. The ICMP discovery might miss some active hosts on the network due to security measures that are put in place to block ICMP attacks.

You configure the timeout value and the number of attempts in the *Discovery Manager* wizard. The ICMP discovery method returns the following information for each detected host:

- IP address: The IP address of the host.
- MAC address: The discovery returns the MAC address only if the Grid member running the discovery is on the same discovered network.

To use the ICMP discovery method, the ICMP protocol between the Grid member performing the discovery and the target networks must be unfiltered.

## NetBIOS

The NetBIOS method queries IP addresses for an existing NetBIOS service. This method detects active hosts by sending NetBIOS queries and listening for NetBIOS replies. It is a fast discovery that focuses on Microsoft hosts or non-Microsoft hosts that run NetBIOS services.

You configure the timeout value and the number of attempts in the *Discovery Manager* wizard. This method returns the following information for each detected host:

- IP address: The IP address of the host.
- NetBIOS name: This value is set to the name returned in the NetBIOS reply.

To use the NetBIOS discovery method, ports 137 (UDP/TCP) and 139 (UDP/TCP) between the Grid member performing the discovery and the target networks must be unfiltered.

## TCP

The TCP discovery probes each active host on a list of TCP ports using TCP SYN packets. This method detects all active hosts that generate SYN ACK responses to at least one TCP SYN. The discovery can determine the OS on a host by analyzing how the host reacts to the requests on opened and closed ports. It then uses the TCP fingerprints to guess the OS. To obtain a TCP fingerprint, IP discovery provides two scanning techniques, SYN and CONNECT.

When you use the SYN technique, the discovery sends a TCP SYN packet to establish a connection on a TCP port. If the port is open, the host replies with a SYN ACK response. The discovery does not close the port connection.

The CONNECT technique is a three-way TCP handshake. The discovery starts with the same process as the SYN technique by sending the TCP SYN packet. If the host replies with a SYN ACK response, the discovery then sends a RST packet to close the connection. If the response contains a RST flag, it indicates that the port is closed. If there is no reply, the port is considered as filtered. The TCP discovery is a deliberate and accurate discovery method. It can basically detect all active hosts on a network provided that there are no firewalls implemented on the network.

You can select the TCP ports, the TCP scanning technique, and configure the timeout value and the number of attempts in the *Discovery Manager* wizard. This method returns the following information for each detected host:

- IP address: The IP address of the host.
- MAC address: The discovery returns the MAC address only if the Grid member running the discovery is on the same discovered network.
- OS: This is set to the highest probable OS reported in the response.

To use the TCP discovery method, the TCP port and a specific set of ports between the Grid member and the discovered networks must be unfiltered. The default set of ports is defined by the factory settings.

## Full

The full discovery method is a combination of an ICMP discovery, a NetBIOS discovery, a TCP discovery, and a UDP scan. This method starts by sending an ICMP echo request. If no IP address on the network responds to the ICMP request, the discovery ends. If there is at least one response to the ICMP echo request, a NetBIOS discovery starts. A TCP discovery then follows by skipping through the active hosts that the NetBIOS discovery detects. The TCP discovery also handles the NetBIOS-detected hosts that have no MAC addresses. This method also performs a UDP scan to determine which UDP ports are open.

You configure the timeout value and the number of attempts in the *Discovery Manager* wizard. The full discovery method returns the following information for each detected host:

- IP address
- MAC address
- OS
- NetBIOS name

To use the full discovery, all the filter and firewall requirements in the ICMP, NetBIOS, and TCP discovery methods apply. The following is a summary of the supported IP discovery methods:

Discovery Type	Returned Data	Guideline	Mechanism
ICMP	<ul style="list-style-type: none"><li>• IP address</li><li>• MAC address</li></ul>	Use ICMP for a rough and fast discovery	ICMP echo request and reply
NetBIOS	<ul style="list-style-type: none"><li>• IP address</li><li>• NetBIOS name</li></ul>	Use NetBIOS for discovering Microsoft networks or non-Microsoft networks that run some NetBIOS services	NetBIOS query and reply
TCP	<ul style="list-style-type: none"><li>• IP address</li><li>• MAC address</li><li>• OS</li></ul>	Use TCP for an accurate but slow discovery	TCP SYN packet and SYN ACK packet
Full	<ul style="list-style-type: none"><li>• IP address</li><li>• MAC address</li><li>• OS</li><li>• NetBIOS name</li></ul>	Use Full for a general and comprehensive discovery	<ol style="list-style-type: none"><li>1. ICMP echo request and reply</li><li>2. NetBIOS query and reply</li><li>3. TCP SYN packet and SYN ACK packet</li></ol>

The method you select to run an IP discovery determines the kind of information the discovery returns and the time it takes to complete an IP discovery. If time is a concern, the following are factors you may consider when configuring an IP discovery:

- The timeout value
- The number of attempts
- The number of ports the discovery scans
- The size of network you want to discover

## Guidelines Before Starting a Discovery

Consider the following guidelines before you start a discovery.

### Database Updates

After the Grid Master receives discovery data from the Grid member, it integrates the data based on the following rules:

- For a discovered host with a new IP address, the appliance marks the IP address "unmanaged."
- For a discovered host associated with one of the following, the appliance updates the data of the associated object:
  - A fixed address reservation or host address reservation
  - A host address not configured for DHCP services
  - A fixed address or host address with the same MAC address as that of the discovered host
  - An A or PTR record
  - A DHCP lease with the same MAC address as that of the discovered host
- For a DHCP lease that does not have any associated object, such as a fixed address or host record, the appliance updates the IP address with the discovered data. When the lease expires and the IP address has no associated objects, the appliance marks the IP address "unmanaged". When the lease expires and the IP address is associated with the same MAC address, the appliance preserves the discovered data.
- For a discovered host associated with one of the following, the appliance updates all data except the MAC address and marks the IP address as a conflict. For information about resolving conflicting addresses, see [Managing Discovered Data](#).
  - A fixed address with a different MAC address than that of the discovered host
  - A DHCP lease with associated objects and with a different MAC address than that of the discovered host
  - An Infoblox host address configured for DHCP services and with a different MAC address than that of the discovered host
- For a discovered host that is part of a DHCP range but does not have a fixed or leased address or is not within an exclusion range, the appliance assigns a DHCP range conflict to the IP address.
- For a discovered host through a vDiscovery, the appliance adds the discovered data to the database. The data is displayed in the IP Map and IP List panels, the **Discovered Data** tab of an object editor, and the Discovered Data section of the IP Address panel.
- The OS of an IP address obtained by an IP discovery supersedes that obtained by a vDiscovery, and the newly discovered name of a host supersedes the last discovered data.
- When a vDiscovery cannot obtain the IP address of a virtual entity, it does not return any discovered data for the entity.
- Only the objects with IP addresses to which the administrators have read/write permission are updated to include the vDiscovery data.

### Database Capacity

When the Grid Master database reaches its maximum capacity (the maximum capacity varies based on the appliance model), the Grid Master stops updating the database and requests that the Grid member stop the discovery. When the discovering Grid member database reaches its capacity, the Grid member pauses the discovery. The appliance displays a dialog to inform you that the discovery pauses. The Grid member resumes the discovery once the database falls below its capacity. When a discovery pauses because of capacity issues, you cannot resume the discovery or start a new discovery. You can check the capacity of your appliance database before starting a discovery.

### HA Failover

In an HA pair, if the Grid Master fails over to the passive node, the passive node takes over and continues with the discovery from the last known state. If an independent appliance fails, the appliance stops the discovery process and keeps the discovery in a paused state. The appliance resumes the discovery once it starts up again.



## Configuring IP Discovery

You must have read/write permission to Network Discovery to initiate a discovery. After you start a discovery, you cannot change the configuration of the discovery, but you can start the discovery process immediately or schedule it for a later date. You can also configure a recurring discovery that repeats on a regular basis. For information, see [Configuring and Starting an IP discovery](#) and [Scheduling IP Discovery](#) below. The appliance saves the configuration of the last discovery. When you start an IP discovery from the IPAM Home, Net Map or Network List panel, you can select the networks on which you want the discovery to run. When you start an IP discovery from the IP Map or IP List panel, the discovered network is the one to which the IP addresses belong. You can include additional networks when you configure the IP discovery from the *Discovery Manager* wizard. You can run an IP discovery on multiple networks in one network view.

### Guidelines for Starting and Scheduling IP Discovery

After you configure a discovery, you can start the discovery process immediately or schedule it for a later date. You can also configure a recurring discovery that repeats on a regular basis. When you start a discovery immediately or schedule for a later date after you configure it, the discovery happens only once and it will not be repeated. To repeat a discovery regularly, you can configure a recurring discovery. A recurring discovery occurs repeatedly based on the schedule you have configured. For more information about how to start a discovery immediately or schedule it for a later date, see [Configuring and Starting an IP discovery](#) and [Scheduling IP Discovery](#) below.

You can configure IP discovery tasks independent of each other and each one contains a specific set of networks and discovery settings.

Note the following guidelines about immediate, regular and recurring IP discovery tasks:

- You cannot run regular and recurring discovery processes concurrently.
- If a recurring discovery is scheduled to start when a discovery is in progress, the recurring discovery will be postponed to the next schedule time. The current recurring discovery will not be performed.
- You can pause and resume all discovery tasks.
- You cannot start a discovery when another one has been paused.
- You cannot use the start command to start a recurring discovery.
- Discovery permissions are applicable to all discovery tasks.
- When you start an IP discovery, only the available IP addresses in the network are discovered. The discovered data through a specific Discovery job can only be modified or deleted by the same Discovery job. When you start a different Discovery job to discover IP addresses in the same network and if some of the information in the network is changed, the newly discovered data and the originally discovered data (now old data) co-exist in the database. When you run the original Discovery job again, the old data is deleted or modified, or the new data is added depending on the information discovered.

### Configuring and Starting an IP discovery

To start an IP discovery immediately after you configure it:

1. From the **Data Management** tab, select the **IPAM** tab, and then select **Discovery** -> **Discover Now** from the Toolbar.
2. In the *Discovery Manager* wizard, click the **General** tab, and then complete the following in the **Basic** tab:
  - **Current Status:** Displays the last discovery status and timestamp. This data is read-only.
  - **Member Name:** Click **Select Member**. In the *Member Selector* dialog box, select the Grid member from which you want to run the discovery. You can also use filters or the Go to function to find a specific member. For information, see [Using Filters](#) and [Using the Go To Function](#).
  - **Merge the discovered data with existing data:** When you select this checkbox, the appliance merges the discovered data with the existing data. It appends newly discovered data to existing data and preserves the existing data when there is no newly discovered data. This checkbox is selected by default. Note if you clear this checkbox, the appliance replaces the existing data with the newly discovered data and if there are no newly discovered values for some fields, the appliance removes the existing values for these fields and the fields become empty.
  - **Update discovered data for managed objects:** Select this checkbox if you want the appliance to update the data of existing managed objects such as A records, PTR records, host records, and fixed addresses,

with the discovered data. If you clear this checkbox, the appliance updates only the unmanaged objects. This checkbox is selected by default.

3. Click **Save** to save the discovery configuration. Note that you must save the configuration before you can start a discovery.
4. Click **Start** to start the IP discovery. You can also do one of the following:
  - **Restore to Defaults:** Restore the discovery configuration using the default values.
  - **Pause:** Stop a running discovery.
  - **Resume:** Resume the discovery that has been stopped.
  - **Save:** Save the discovery configuration.
  - **Close:** Cancel the configuration. If you have started a discovery, the discovery runs in the background when you click **Close**. For information, see [Running Tasks in the Background](#).



#### Note

Once you start a discovery, you cannot change the discovery configuration. After you click **Start**, the button changes to **Pause**. You can click **Pause** to pause a discovery. When the discovery is paused, the button changes to **Resume**. You can click **Resume** to continue the paused discovery.

## Defining IP Discovery Method

To configure the IP discovery method you want to use:

1. From the **Data Management** tab, select the **IPAM** tab, and then select **Discovery** -> **Discover Now** from the Toolbar.
2. In the *Discovery Manager* wizard, click the **IPv4 Device Discovery** tab, and then complete the following in the **Basic** tab:
  - **Mode:** Select the IP discovery method you want to use. For information, see [Supported IP Discovery Methods](#). If you select **TCP** or **FULL**, ensure that you configure the TCP ports in the **Advanced** tab. The default is **FULL**.
  - Click the Add icon to add networks. In the *Network Selector* dialog box, select the network view and networks. Use SHIFT+click and CTRL+click to select multiple networks. You can also use filters or the Go to function to find a specific network. For information, see [Using Filters](#) and [Using the Go To Function](#). You can do the following in the table:
    - Click the Add icon again to add more networks.
    - Select a network or multiple networks in the network table and click **Delete** to delete them.
    - Click the Export icon to export the data in CSV format.
    - Click the Print icon to print the data.
  - **Disable:** Select this to exclude an IP discovery task. IP discovery is enabled by default.
3. If you select **TCP** or **FULL** in **Mode**, click the **Advanced** tab and complete the following:
  - **TCP Scan Technique:** Select the TCP technique you want to use for the discovery. The default is SYN. For information, see [TCP](#).
  - In the port table, select the checkbox of the port you want to configure. You can select all ports by clicking the checkbox in the header. Optionally, you can click the Add icon and complete the following to add a new service to the list.
    - Port: Enter the port number you want to add to the list. You must enter a number between 1 and 65535.
    - Service: Enter the name of the service.You can also delete a specific TCP port in the list. You can select multiple ports for deletion.
  - **Timeout (ms):** Enter the timeout value in milliseconds for the discovery. The timeout value determines how long the discovery waits for a response from an IP address after probing it. The minimum is 5 and the maximum is 4000. The default is 1000.
  - **Attempts:** Enter the number of times you want the discovery to probe an IP address when scanning a network. The minimum is 1 and the maximum is 5. The default is 2.
4. Click **Save** to save the discovery configuration. Note that you must save the configuration before you can start a discovery.
5. Click **Start** to start the IP discovery. You can also do one of the following:

- **Restore to Defaults:** Restore the discovery configuration using the default values.
- **Pause:** Pause a running discovery.
- **Resume:** Resume the discovery that has been stopped.
- **Save:** Save the discovery configuration.
- **Close:** Cancel the configuration. If you have started a discovery, the discovery runs in the background when you click **Close**. For information, see [Running Tasks in the Background](#).

You can do the following after a discovery is complete:

- View the discovery status. You can view the current discovery status in the *Discovery Status* widget on the Dashboard. For information, see [Dashboards](#).
- View the discovered data. For information, see [Viewing Discovered Data](#).
- Manage the discovered data. For information, see [Managing Discovered Data](#).

## Scheduling IP Discovery

After you configure a discovery (as described in [Configuring and Starting an IP discovery](#) above), you can schedule to run a one-time IP discovery at a later date and time. You can also schedule a recurring IP discovery by configuring a recurrence pattern. The appliance automatically starts a recurring discovery based on the configured schedule and detects any newly added or removed networks. Note that you can only schedule the start of a discovery, you cannot schedule it to pause, stop, or resume. After a scheduled discovery starts, you can then pause, stop, or resume it. You can schedule only one IP discovery at a time. Once you schedule a discovery, you cannot change the configuration until the task is cancelled or executed. You can however disable a recurring discovery. When you disable a recurring discovery, it will not recur during the scheduled interval.

To schedule a one-time IP discovery for a specific date and time:

1. From the **Data Management** tab, select the **IPAM** tab, and then select **Discovery -> Discover Now** from the Toolbar.  
or  
From the *Discovery Status* widget, select **Discover Now** from the drop-down list, and then click **Discovery Manager**.
2. In the *Discovery Manager* wizard, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, click **Later**, and then specify a date, time, and time zone. Click **Schedule Start** to schedule the discovery. If applicable, you can select **Click here to view/manage the scheduled items to reschedule a discovery or view all scheduled discoveries**.

To schedule a recurring IP discovery:

1. From the **Data Management** tab, select the **IPAM** tab, and then select **Discovery -> Schedule Discovery** from the Toolbar.  
or  
From the *Discovery Status* widget, select **Schedule Discovery** from the drop-down list, and then click **Discovery Manager**.
2. In the *Discovery Manager* wizard, complete the following in the scheduler:
  - **Disable:** Select this checkbox to disable recurring IP discovery. When you select this checkbox, IP discovery will not recur during the scheduled interval. Clear the checkbox to enable recurring IP discovery.
 If you select **Hourly**, complete the following:
  - a. **Schedule every hour(s) at:** Enter the number of hours between each update instance. You can enter a value from 1 to 24.
  - b. **Minutes past the hour:** Enter the number of minutes past the hour. For example, enter 5 if you want to schedule the rule update five minutes after the hour.
  - c. **Time Zone:** Select the time zone for the scheduled time from the drop-down list.

If you select **Daily**, you can select either **Everyday** or **Every Weekday** and then complete the following:

- **Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
- **Time Zone:** Select the time zone for the scheduled time from the drop-down list.

If you select **Weekly**, complete the following:

- **Schedule every week on:** Select any day of the week.
- **Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
- **Time Zone:** Select the time zone for the scheduled time from the drop-down list.

If you select **Monthly**, complete the following:

- **Schedule the day of the month:** Enter the day of the month and the monthly interval. For example, to schedule the rule update on the first day after every 2 months, you can enter Day 1 every 2 month(s).
- **Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
- **Time Zone:** Select the time zone for the scheduled time from the drop-down list.

3. Click **Save** to save the configuration.

If the discovery task fails during a scheduled interval, the task stops and will not continue for the corresponding occurrence. The scheduled task resets and the discovery starts at the next scheduled time. For example, when you configure a recurring discovery to occur every five hours, discovery starts at the following hours on each day: 00:00, 05:00, 10:00, 15:00, and 20:00. If the discovery scheduled for 05:00 fails, the discovery starts at the next recurrence, which is at 10:00. For information about failed discovery, see [Guidelines for Starting and Scheduling IP Discovery](#). The following examples explain when a recurring discovery starts based on your configuration:

#### Example 1

When you configure a recurring discovery to occur every five hours, the discovery starts at the following hours on each day: 00:00, 05:00, 10:00, 15:00, and 20:00. The first occurrence on each day starts at 00:00.

#### Example 2

When you configure a recurring discovery to occur every two days during a week, the discovery starts on the following days every week: Monday, Wednesday, Friday, and Sunday. The first occurrence starts on Monday of each week.

## Monitoring IP Discovery Status

You can monitor IP discovery status through the *Discovery Status* widget on the Dashboard. You can also start, pause, resume, and stop a discovery from the widget. For information, see [Discovery Status](#).

## Configuring vDiscovery Jobs

A vDiscovery job retrieves information about virtual entities in cloud environments that are managed through cloud management platforms (CMPs) such as VMware, OpenStack, AWS, Azure, and GCP. The current vDiscovery feature supports tenants, networks, and compute VMs only. It does not support data that is retrieved from load balancer networks, load balancer VMs, Kubernetes platform VMs, application gateways, service VMs, SQL VMs, or any other VMs that are created using cloud services such as Kubernetes service or analytics service, where the IPAM is handled by the respective orchestration engines of the cloud provider.

Note that if the vDiscovery job retrieves unsupported data from AWS, Azure, or GCP, then it impacts the performance of the vDiscovery process.

You must first select a member to run the vDiscovery job. To ensure that the job is executed properly, verify the connection between the discovering member and the discovered endpoint. If you select HTTPS as the protocol for communication, you must upload either an SSL CA (Certified Authority) certificate or a self-signed SSL certificate to the Grid. If the certificate has expired, ensure that you delete the expired certificate and then upload a new certificate. For more information about uploading certificates, see [Managing Certificates](#).



#### Note

After you upload a new certificate, wait for two minutes before running a vDiscovery job. This is because the newly uploaded certificate takes some time to be reflected in the NIOS database.

Note that when you disable any virtual entities or interfaces on the CMP, the appliance excludes them from the vDiscovery job. In situations where the discovering member you select to perform vDiscovery jobs is disconnected from the Grid Master, the member continues to execute vDiscovery jobs based on the configured schedule. Newly discovered

data replaces previously discovered data. The last set of discovered data is considered the most up-to-date and is sent to the Grid Master when the member reconnects with the Grid Master.

When you configure vDiscovery jobs, you can enable the Infoblox NIOS appliance to automatically create DNS records for discovered IP addresses of VM instances that are served by the NIOS appliance. You can configure the appliance to add DNS records for specific DNS views associated with the network view defined for public and private IP addresses of VM instances served by the appliance. For more information about this feature, see [Creating DNS Records for Newly Discovered VMs](#) below.

Before you configure and start a vDiscovery job, there are a few guidelines to consider.

Sections covered in this topic are:

- [Guidelines for Configuring vDiscovery Jobs](#)
- [Creating a new vDiscovery Job](#)
  - [Selecting the vDiscovery Member](#)
  - [Selecting the Endpoint Server](#)
  - [Defining Network Views](#)
  - [Defining Filtering Options \(for AWS and GCP only\)](#)
  - [Defining Policies for Handling Discovered Data](#)
  - [Creating DNS Records for Newly Discovered VMs](#)
  - [Scheduling vDiscovery Jobs](#)
- [Managing vDiscovery Jobs](#)

## Guidelines for Configuring vDiscovery Jobs

Consider the following guidelines before starting a vDiscovery job:

- Discovered data through a specific vDiscovery job can only be modified by the same vDiscovery job. If you create a different vDiscovery job to discover the same network and some of the information in the network has changed, the newly discovered data and the original discovered data (now old data) co-exist in the database. For example, vDiscovery job "AWSJob1" discovers IP address 10.0.0.11 with VM name "corpxyz." This discovered data is stored in the NIOS database. You subsequently create a new vDiscovery Job "AWSJob2" to discover the same IP but now the VM name has been changed to "corp200." The discovered VM name for 10.0.0.11 is now "corp200." Both VM names "corpxyz" and "corp200" exist in the database until you run AWSJob1 again, and "corpxyz" will be removed from the database.
- If the "ERROR: PycURL" error is displayed when you run a vDiscovery job, it is possible that the cloud provider has updated their certificate. You need to download the latest certificate from the cloud provider website and upload it to NIOS. For example, for AWS, download the certificates from <https://www.amazontrust.com/repository/>. For more information see [Error while running job](#) below.
- Cloud Extensible Attributes:
  - To merge discovered cloud extensible attributes into NIOS, you must have at least one cloud license (Cloud Network Automation or Cloud Platform) installed in the Grid. However, if you want to view the discovered cloud data in the **Cloud** tab of Grid Manager, you must have the Cloud Network Automation license installed on the Grid Master. Otherwise, even though the cloud data is merged into the NIOS database, you cannot view it through Grid Manager. For more information about Cloud Network Automation, see [Deploying Cloud Automation Network](#).
  - In addition, ensure that you select the auto consolidation options when defining policies for how to handle the discovered cloud extensible attributes, as described in [Defining Policies for Handling Discovered Data](#). Note that only cloud extensible attributes for managed objects are updated. To update cloud extensible attributes for unmanaged objects, you can first convert the unmanaged objects to managed objects. For more information about [Managing Unmanaged Data](#), see [Managing Discovered Data](#).
- When you select **VMware** as the endpoint server for the vDiscovery job, consider the following:
  - Typically, vDiscovery does not collect data about "Network" from VMware vSphere and vCenter servers. Therefore, you must first define a network in NIOS in order to discover IPs in the network.
  - However, if you use vCenter to define a network as "IP\_Pool," which contains the CIDR for the network, vDiscovery is able to collect this data and translate it into a network. NIOS then creates the network and updates it with corresponding Cloud extensible attributes.

- If you use an AWS Elastic IP address and select **AWS** as the endpoint server for the vDiscovery job, you must first manually create the network in NIOS before launching the vDiscovery job in order to discover the Elastic IP.
- Properties of a VM address (such as interface name, network encapsulation, segmentation, VLAN ID) can only be updated through cloud extensible attributes, depending on the policies you select when configuring a vDiscovery task. To properly integrate discovered VM address properties with NIOS, ensure that you do one of the following:
  - Define cloud extensible attributes for these properties through your cloud adapter. For information about how to define cloud extensible attributes, refer to the *Quick Start Guide* for your cloud adapter, available on the Support site.
  - Convert unmanaged VM addresses to managed so discovered cloud extensible attributes can be merged into NIOS. For more information, see [Converting Unmanaged Data](#).
  - Unmanaged objects in NIOS are updated with newly discovered data and stay as unmanaged objects. Managed objects in NIOS that have no conflict with newly discovered data are updated and stay as managed objects.
  - If there is conflict between managed objects in NIOS and newly discovered data, the managed objects are not updated and stay as managed objects while a conflict is flagged for these objects. For more information about how to resolve the conflict, see [Resolving Conflicting Addresses](#).
  - NIOS does not automatically remove unmanaged objects that were discovered in the past but do not exist in current discoveries. This can happen if your network topology has changed in between discoveries. You can manually remove these unmanaged objects if you do not want them to stay in the database. For more information about how to remove these objects, see [Clearing Unmanaged Data](#) and [Clearing Discovered Data](#).
  - For delegated DNS and DHCP objects, changes are handled by the delegated Cloud Platform Appliance based on the scope of delegation. Discovered data is still updated by the Grid Master.
- You cannot delete discovered tenants, networks, and VMs through the vDiscovery process. Conflict management is not supported for these objects. Note the following:
  - Discovered subnets are always created as managed networks. Pools of IP addresses (public pools) when discovered are translated into managed networks. Any discovered public IP that is not linked to a pool is marked as unmanaged data in NIOS, unless there is a corresponding network already created.
  - Tenant and VM information is merged into NIOS through cloud extensible attributes, only if there is a cloud license (Cloud Network Automation or Cloud Platform) installed in the Grid. Properties for unmanaged tenants and VMs are always updated while properties for managed tenants and VMs are updated only if auto-consolidation of tenant and VM information is selected when you configure the vDiscovery job.
- If a newly created vDiscovery job fails, first ensure that no control characters or permanently undefined Unicode characters such as non-breaking space, non-breaking hyphen, and so on are present in the job name.

## Creating a new vDiscovery Job

To create a new vDiscovery job, complete the following tasks:

1. Name the new job and select a member to perform the vDiscovery, as described in [Selecting the vDiscovery Member](#) below. Do not use control characters or undefined Unicode characters in your job name.
2. Select a cloud platform for the vDiscovery, as described in [Selecting the Endpoint Server](#) below.
3. Define network views for public and private IP addresses, as described in [Defining Network Views](#) below.
4. Define policies for handling discovered data, as described in [Defining Policies for Handling Discovered Data](#) below.
5. Optionally, you can enable NIOS to automatically create DNS records for newly discovered VMs using their IP addresses, as described in [Creating DNS Records for Newly Discovered VMs](#) below.
6. Schedule the vDiscovery job, as described in [Scheduling vDiscovery Jobs](#) below.

## Selecting the vDiscovery Member

To create a new or modify an existing vDiscovery job:

1. For a new vDiscovery job: From the **Data Management** tab, select the **IPAM** tab, then select **vDiscovery -> New** from the Toolbar; or from the **Cloud** tab, select **vDiscovery -> New** from the Toolbar.  
or  
To modify an existing job: From the **Data Management** tab, select the **IPAM** tab and click **vDiscovery -> Discovery**



**Manager** from the Toolbar, or from the **Cloud** tab, select **vDiscovery** -> **Discovery Manager** from the Toolbar. In the **vDiscovery Job Manager** editor, click the Action icon

next to a selected job and select **Edit** from the menu.

- In step one of the *vDiscovery Job* wizard or in the **General** tab of the *vDiscovery Job Properties* editor, complete the following:
  - Job Name:** Enter the job name for this vDiscovery. It might be helpful to use a name that is unique to this specific discovery if you plan to configure multiple vDiscovery jobs. You cannot update the job name after the vDiscovery job is run for the first time.
  - Member:** Click **Select** to choose the Grid member that will perform the vDiscovery job. If only a single member is active, the appliance name automatically appears here. When you select a Cloud Platform Appliance to perform vDiscovery, it communicates directly with the CMPs to obtain information that is not available through the provisioning process from the cloud adapter.
  - Comment:** Enter information to describe this discovery.  
The new job will not execute until you have completed all configuration steps in the wizard. You will not be able to save this job until you have completed all job settings.
- Click **Next** to select an endpoint server on which you want to perform the vDiscovery job, as described in *Selecting the Endpoint Server* below, or save the configuration after you have modified data in this tab.

### Selecting the Endpoint Server



#### Note

- You might lose some discovered data if you modify any of the following parameters for an existing vDiscovery job. To avoid this, create a new vDiscovery job instead.
- Azure Government Cloud uses different service endpoints for its services. For more information about Azure service endpoints, refer to the *Infoblox Installation Guide for vNIOS for Azure*.
- Updates to the root CAs of Azure services installed by Microsoft, can cause vDiscovery to fail. If vDiscovery fails with ERROR: PycURL error: (60, 'SSL certificate problem: unable to get local issuer certificate'):
  - Download the **DigiCert Global Root G2** Certificate from [DigiCert Root Certificates](#).
  - Upload the certificate to NIOS as described in [Uploading CA Certificates](#).
- From NIOS 8.4.2 onward, each GCP vDiscovery job can utilize only one uniquely named service account file.
- You can use the same service account file for different GCP vDiscovery jobs in earlier NIOS releases (such as 8.4.0, 8.4.1). However, deleting one vDiscovery job causes the other vDiscovery job with the same name to be stuck in "Job in Progress" state. It is recommended to use only one uniquely named service account file for each vDiscovery job.
- You cannot edit an existing GCP vDiscovery job to upload a different service account file.
- Auto created networks and VM instances having IP address of auto created networks cannot be discovered by GCP vDiscovery.

- In step two of the *vDiscovery Job* wizard, or in the **Endpoint** tab of the *vDiscovery Job Properties* editor, complete the following:
  - Server Type:** Choose one of the following server types for this vDiscovery:
    - AWS:** Collects information available for the AWS service endpoint. You can perform vDiscovery jobs through a proxy server in an AWS deployment, including Amazon Route 53. For more information, see [Configuring Proxy Servers](#).
    - Azure:** Collects information available for virtual entities in the specified VNets (Azure virtual networks) within the Microsoft Cloud.
    - OpenStack:** When you select this server type, vDiscovery discovers network information stored in Neutron servers, VM instance information in Nova servers, and tenant or project information in Keystone servers.
    - VMware:** Supports VMware vCenter and vSphere servers v5.0 and later. Collects information for all virtual entities running on the specified servers.
    - GCP:** Collects information available for virtual entities in the specified GCP project.

Depending on the server type you select, other options in this step change accordingly, as follows:

For **AWS**, complete the following:

- **Service Endpoint:** This is typically the regional service endpoint for the desired Amazon region. Example: `ec2.us-west-1.amazonaws.com`. For more information about AWS service endpoints, refer to the *Infoblox Installation Guide for vNIOS for AWS*, available on the Infoblox Support site. For a list of available AWS service endpoints, see <https://docs.aws.amazon.com/general/latest/gr/ec2-service.html>
- **Port:** Enter the port you want to use for the vDiscovery job.
- **Protocol:** The protocol used for AWS is always over SSL. AWS provides certificates that is linked to the CA. By default, this certificate is embedded in NIOS and used as a reference for the CA when connecting to AWS. You can also upload a new certificate as described in [Managing Certificates](#). If you upload a new certificate, the embedded certificate will be overwritten by the new one.
- **Allow unsecured connection:** This option is not applicable for AWS connection.

**Credentials:** Select the method you want to use to authenticate the connection between the Grid member and AWS for discovery jobs. You can select one of the following:

- **Use instance profile:** An instance profile is a container for an IAM role that you use to pass role information to an EC2 instance when the instance is up and running. Select this option if you want to collect information from AWS by waiving a user's credentials and by using configuration of a predefined IAM role to get a temporary token that allows API calls. Note that you must first configure the option for "instance profile" in AWS, define an IAM role in the instance profile, and then set permissions for this role before you can select this option. Otherwise, this option is disabled. When you select this, you do not need to provide user credentials.  
**Note:** Multiple vDiscovery jobs are required for multiple accounts irrespective of the scope of the instance profile.
- **Use IAM credential:** Select this if you want to authenticate by using IAM roles to grant secure access to AWS resources from your EC2 instances. Click **Select** to choose the IAM role and use its credentials to access AWS resources from your EC2 instances when they are up and running.
  - **Access Key ID and Secret Access Key:** Enter the Access Key ID and Secret Access Key for the AWS service endpoint. This is the secret key pair for the administrator account that executes the discovery job. For more information, refer to the *Infoblox Installation Guide for vNIOS for AWS*, available on the Infoblox Support site.

For more information about instance profiles and IAM roles, refer to the AWS documentation.

For **Azure**, complete the following:

- **Service Endpoint:** This is the service endpoint for the desired VNet in the Microsoft Cloud. For more information about Azure service endpoints, refer to the *Infoblox Installation Guide for vNIOS for Azure*.
- **Port:** Enter the port you want to use for the vDiscovery job.
- **Protocol:** The protocol used for Azure is always over SSL. Azure provides certificates that is linked to the CA. By default, this certificate is embedded in NIOS and used as a reference for the CA when connecting to Azure. You can also upload a new certificate as described in [Managing Certificates](#). If you upload a new certificate, the embedded certificate will be overwritten by the new one.
- **Allow unsecured connection:** This option is not applicable for Azure connection.
- **Client ID and Client Secret:** Enter the client ID and client secret for the Microsoft Azure account. When you configure the client account, ensure that you have authorization to obtain device information on a wide network basis. If you replace the client secret of the vDiscovery job with the existing client ID, you must restart the vDiscovery job for the changes to take effect. For information about Azure client ID and client secret, refer to Microsoft Azure documentation.

For **OpenStack**, complete the following:

- **Keystone Server IP:** Enter the Keystone Server IP address.
- **Keystone Server Port:** Select this checkbox if you are using Identity Endpoint for the connection, else enter the Keystone Server Port number.
- **Protocol:** Select **HTTP** or **HTTPS** as the protocol. When you select **HTTPS**, you must upload the corresponding SSL CA certificate to the Grid in order for NIOS to communicate with OpenStack, as described in [Managing Certificates](#).
- **Allow unsecured connection:** This option is enabled when you use HTTPS as the protocol. When you select this, the appliance bypasses remote SSL certificate validation. Select this option only if security for the HTTPS



connection between the discovering member and OpenStack is irrelevant, or if the connection is protected by other security measure besides TSL/SSL, such as an isolated private circuit.

- **Username** and **Password**: Enter the username and password of the administrative account that was configured on OpenStack. When you configure the administrative account, ensure that you have authorization to obtain device information on a wide network basis.
- **Identity Version**: Select the Keystone server identity service version from the drop-down list. You can select one of the following: **Keystone v2** and **Keystone v3**. By default, **Keystone v2** is selected.
- **Domain Name**: Enter the domain name. This field is displayed only if you select **Keystone v3** as the identity version.

For **VMware**, complete the following:

- **Host**: Enter the host name of the VMware server.
- **Port**: Enter the port number of the VMware server.
- **Protocol**: Select **HTTP** or **HTTPS** as the protocol. When you select **HTTPS**, you must upload the corresponding SSL CA certificate in order for NIOS to communicate with the VMware server, as described in [Managing Certificates](#).
- **Allow unsecured connection**: This option is enabled when you use HTTPS as the protocol. When you select this, the appliance bypasses remote SSL certificate validation. Select this option only if security for the HTTPS connection between the discovering member and VMware is irrelevant, or if the connection is protected by other security measure besides TSL/SSL, such as an isolated private circuit.
- **Username** and **Password**: Enter the username and password of the administrative account that was configured on the specified VMware server. When you configure the administrative account, ensure that you have authorization to obtain device information.

For **GCP**, complete the following:


- **Service Account File**: Click **Upload** to select and upload the GCP service account file from your local disk. For more information about creating a GCP Service Account, see [Creating GCP Service Account](#).

Click **Next** to define the network views to which discovered data belongs for both public and private IP addresses, as described in Defining NetworkViews below.

## Defining Network Views

To track overlapping networks and IP address ranges so you can discover specific networks and IP addresses, you can associate one or more network views with a Grid member or Cloud Platform Appliance that is selected to run the vDiscovery. You can then define a specific network view to which discovered data for public and private IP addresses belongs if the network view is not automatically detected. If no network view is specified, the default network view is used.

For Network Insight, when you discover networks using multiple discovery interfaces, you must configure network views so you can associate each discovery interface with an available network view. Note that on the same discovering member, each discovery interface must have a unique network view association.

1. For a new vDiscovery job: From the **Data Management** tab, select the **IPAM** tab, then select **vDiscovery** -> **New** from the Toolbar; or from the **Cloud** tab, select **vDiscovery** -> **New** from the Toolbar.  
or  
To modify an existing job: From the **Data Management** tab, select the **IPAM** tab and click **vDiscovery** -> **Discovery Manager** from the Toolbar, or from the **Cloud** tab, select **vDiscovery** -> **Discovery Manager** from the Toolbar. In the **vDiscovery Job Manager** editor, click the Action icon next to a selected job, and then select **Edit** from the menu.  
*Action icon*  

2. In step three of the *vDiscovery Job* wizard, or on the **Network View** tab of the *vDiscovery Job Properties* editor, complete the following:
  - Under the **For Public IP Addresses** section, select one of the following options the appliance uses if the network view is not automatically detected:
    - **This Network View**: From the drop-down list, specify a network view to which discovered data for public IP addresses belongs. The default is the default network view. You cannot create a new network view for this option.

- **The tenant's network view (if it does not exist, create a new one):** Select this only if at least one cloud license is installed in the Grid. When you select this, discovered data for public IP addresses belongs to the tenant's network view. If the network view does not exist, the appliance creates it (only if a cloud license is installed in the Grid). The appliance uses tenant information of a discovered public IP address to create a new NIOS network view for all discovered objects (primarily subnets) for that tenant. For example, AWS tenants by default are associated with the user account's 12-digit account number (such as 2233441247523), which is the identifier for all objects that are created by that account in AWS. That tenant value becomes the identifier for the new network view as its objects are discovered by NIOS.
  - Under the **For Private IP Addresses** section, select one of the following options the appliance uses if the network view is not automatically detected:
    - **This Network View:** From the drop-down list, select a network view to specify a network view to which discovered data for private IP addresses belongs. The default is the default network view. You cannot create a new network view for this option.
    - **The tenant's network view (if it does not exist, create a new one):** Select this only if at least one Cloud Platform Appliance is active or a cloud license is installed in the Grid. When you select this, discovered data for private IP addresses belongs to the tenant's network view. If the network view does not exist, the appliance creates it (only if a cloud license is installed in the Grid). The appliance uses tenant information of a discovered private IP address to create a new NIOS network view for all discovered objects (primarily subnets) for that tenant. For example, AWS tenants by default are associated with the user account's 12-digit account number (such as 2233441247523), which is the identifier for all objects that are created by that account in AWS. That tenant value becomes the identifier for the new network view as its objects are discovered by NIOS.
3. Configure filtering options for vDiscovery as described in the *Defining Filtering Options* section.

#### Defining Filtering Options (for AWS and GCP only)

Configure a vDiscovery job with CIDR-based filters to restrict vDiscovery to discover data only in included networks or skip vDiscovery in excluded networks. The maximum number of networks you can include in or exclude from a vDiscovery job is 10.



#### Note

You can enable the vDiscovery filter options only for private networks, not for public networks.

To define the filtering options for a vDiscovery job, complete the following steps:

1. Do one of the following:
  - To create a new vDiscovery job:
    - On the **Data Management** tab -> **IPAM** tab, from the Toolbar, select **vDiscovery** -> **New**.  
or
    - On the **Cloud** tab, from the Toolbar, select **vDiscovery** -> **New**.
  - To modify an existing job:
    - i. Do one of the following:
      - On the **Data Management** tab -> **IPAM** tab, from the Toolbar, select **vDiscovery** -> **Discovery Manager**.  
or
      - On the **Cloud** tab, from the Toolbar, select **vDiscovery** -> **Discovery Manager**.
    - ii. In the *vDiscovery Job Manager* editor, click the **Action** icon next to a selected job, and then select **Edit** from the menu.
2. In step three of the *vDiscovery Job* wizard, do the following in the **Filtering Options** section:
  - a. Select **Enable Filter**.
  - b. In the **Network** drop-down list, do one of the following:
    - Select **INCLUDE** to include networks in which the vDiscovery job must discover resources. All other networks will be excluded from vDiscovery.
    - Select **EXCLUDE** to specify networks that must be excluded from vDiscovery. All other networks will be included in vDiscovery.

- c. Click the **Add** icon and, for networks to be included in or excluded from vDiscovery, enter a valid IP address range in the CIDR notation.
3. Click **Next** to configure the appliance to handle discovered data as described in the *Defining Policies for Handling Discovered Data* section.

## Defining Policies for Handling Discovered Data

In this step, you define how the appliance handles discovered data.

1. For a new vDiscovery job: From the **Data Management** tab, select the **IPAM** tab, then select **vDiscovery** -> **New** from the Toolbar; or from the **Cloud** tab, select **vDiscovery** -> **New** from the Toolbar.  
or  
To modify an existing job: From the **Data Management** tab, select the **IPAM** tab and click **vDiscovery** -> **Discovery Manager** from the Toolbar, or from the **Cloud** tab, select **vDiscovery** -> **Discovery Manager** from the Toolbar. In the **vDiscovery Job Manager** editor, click the Action icon

next to a selected job and select **Edit** from the menu.

2. In step four of the *vDiscovery Job* wizard, or in the **Data Consolidation** tab of the *vDiscovery Job Properties* editor, complete the following:

Under **When inserting discovered into NIOS**, select one or both of the following:

- **Merge the discovered data with existing data:** When you select this checkbox, the appliance merges the discovered data with the existing data. It appends newly discovered data to existing data and preserves the existing data when there is no newly discovered data. This checkbox is selected by default. If you clear this checkbox, the appliance replaces the existing data with the newly discovered data and if there are no newly discovered values for some fields, the appliance removes the existing values for these fields and the fields become empty.
- **Update discovered data for managed objects:** Select this checkbox if you want the appliance to update discovered data for all corresponding NIOS objects (if they exist in NIOS). If you do not select this checkbox, the appliance updates only the discovered data for unmanaged objects. None of the managed data will be updated. This checkbox is selected by default.
- **For every newly discovered IP address, create:** Select this checkbox to enable NIOS to automatically create or update DNS records for discovered network entities and VM instances. It does not include cloud adapters such as AWS or DDNS. This is applicable if the records were originally created by vDiscovery. If you select this checkbox, NIOS considers all records created in a zone as one and calculate it as one serial number change.
  - **Host:** Select this to automatically create Host records for discovered entities.
  - **A & PTR Record:** Select this to automatically create A and PTR records for discovered entities. Note that the DNS zones and reverse-mapping zones to which the records belong must exist in NIOS before the vDiscovery job is executed. Otherwise, vDiscovery does not create the records.
  - **The DNS name will be computed from the formula:** Enter the formula that NIOS uses to create the DNS records for each discovered VM address. For example, if there are two IP addresses associated with a VM, NIOS creates two DNS records, or a host record with two IP addresses, depending on your configuration. You must use the syntax of `${parameter name}` for the formula.

For AWS, Azure, GCP, OpenStack, and VMware cloud platforms, this field supports the following parameters:

```
vm_id, vm_name, discovered_name, tenant_id,
tenant_name, subnet_id, subnet_name, network_id, network_name,
vport_name, ip_address, ip_address_octet1 or 1, ip_address_octet2
or 2, ip_address_octet3 or 3, ip_address_octet4 or 4. Note that it does
not support IPv6 addresses.
```

For example, when you enter `${vm_name}.corpxyz.com` and the discovered `vm_name` = XYZ, the DNS name for this IP becomes

```
XYZ.corpxyz.com
```

. When you enter `${discover_name}` here and the discovered name for the IP is `ip-172-31-1-64.us-west-1.compute.internal`, the DNS name for this IP is `ip-172-31-1-64.us-west-1.compute.internal`.

- Under **Select the DNS view to which the DNS records are being added**, select one or both of the following:
  - **Use this DNS view for public IPs:** Select this checkbox to add DNS records to a specific DNS view for public IPs. Select a DNS view from the drop-down list. If you do not select a DNS view, the DNS records are added to the default DNS view.
  - **Use this DNS view for private IPs:** Select this checkbox to add DNS records to a specific DNS view for private IPs. Select a DNS view from the drop-down list. If you do not select a DNS view, the DNS records are added to the default DNS view. If you are changing the DNS view, ensure that the **Merge the discovered data with existing data** checkbox is not selected. Note that the **Use this DNS view for public IPs** and **Use this DNS view for private IPs** fields will be disabled, if you select **The tenant's network view (if it does not exist, create a new one)** option when you define the network views to which discovered data belongs for both public and private IP addresses, as described in Defining Network Views above. Under **When discovered data is linked to managed data**, select any combination of the following. Tenants and VMs are managed objects when they have NIOS objects, such as host records or fixed addresses, associated with them. Otherwise, they are unmanaged objects. The appliance always updates properties for all unmanaged objects.
  - **Auto-consolidate on managed Tenant's properties:** When you select this checkbox, the appliance updates properties with discovered data for managed tenants, as well as unmanaged tenants (NIOS always updates unmanaged tenants). When you clear this checkbox, the appliance does not update discovered data for managed tenants. This checkbox is selected by default.
  - **Auto-consolidate managed VM's properties:** When you select this checkbox, the appliance updates properties and extensible attributes with discovered data for managed VMs, as well as unmanaged tenants (NIOS always updates unmanaged tenants). When you clear this checkbox, the appliance does not update discovered data for managed VMs. This checkbox is selected by default.
  - **Auto-consolidate Cloud EAs on managed data:** When you select this checkbox, NIOS updates discovered extensible attribute values for managed objects that contain cloud extensible attributes, only if a cloud license is installed in the Grid. This includes the update of the extensible attribute **VM ID** (which links the NIOS object to the VM) whenever a VM is added, updated or removed depending on the information collected. As a result, when a VM instance reuses an IP address or when a VM instance is deleted in the Cloud, the DNS Records or fixed address tied to that IP address are also updated, reflecting the new value of the VM instance ID. To update cloud extensible attributes for unmanaged objects, convert the objects to managed objects in NIOS. For more information, see [Managing Unmanaged Data](#). The extensible attribute **VM ID** is not updated if you do not enable the **Auto-consolidate Cloud EAs on managed data** checkbox, which leads to a conflict on that IP address. The NIOS object does not link to the same VM as the newly discovered IP. In such cases, you can use the **Resolve Conflicts** option to update either your NIOS objects or your discovered data. For information about resolving conflicts, see [Resolving Conflicting Addresses](#).
- 3. Click **Next** to schedule this vDiscovery job and specify when the job should start, as described in Scheduling vDiscovery Jobs below.

### Creating DNS Records for Newly Discovered VMs

When you configure the policies NIOS uses to handle discovered data, you can enable NIOS to automatically create or update DNS records for discovered IP addresses of VM instances. NIOS automatically adds Host records or A and PTR records for the discovered VMs based on your configuration. You can also enter a formula that NIOS uses to create the

DNS name for the discovered IP based on its VM parameters such as VM name or discovered name. By doing so, NIOS is able to discover public and private IP addresses by looking up the corresponding DNS names.

Discovered data includes IP addresses for the VMs and the associated information such as VM ID, VM Name, Tenant ID, and others. Note that corresponding zones must already exist in order for NIOS to add DNS records. Otherwise, NIOS does not add any DNS records and it logs a message in the syslog. For information about how to enable this feature, see Defining Policies for Handling Discovered Data above.

NIOS automatically adds DNS records based on the following conditions:

- The corresponding DNS zones must already exist in the NIOS database. NIOS does not automatically create DNS zones for the records.
- To create a PTR record, the corresponding reverse-mapping zone must exist.
- A DNS zone cannot be associated with more than one DNS view. NIOS does not create DNS records for zones that are associated with multiple DNS views.
- NIOS adds new DNS records only if the VM name for the discovered IP address is available and there is no conflict with information about the associated network view.

On subsequent vDiscovery, if an IP for a VM is removed, the corresponding DNS records are removed. If the IP for a VM is changed, the IP address in the corresponding DNS record is changed accordingly. If the DNS record name template is changed, all the DNS records are replaced with the DNS records using the new template. All administrative actions for these changes are recorded in the Audit log. Summary of the changes are logged in the syslog.

The following table captures some scenarios about how vDiscovery handles various actions and what the outcome is for the information on the Cloud Platform appliance and in the NIOS database. All the scenarios in the table use the following template: \$(discovered\_name).



**Note**

vDiscovery updates only records that are created by the vDiscovery process. It does not create or update DNS records that are originally created by other admin users.

Actions and Conditions	Cloud Platform Data before vDiscovery	Cloud Platform Data after vDiscovery	NIOS Data before vDiscovery	NIOS Data after vDiscovery
<ul style="list-style-type: none"> <li>• Add new VM (vma) on Cloud Platform appliance</li> <li>• Automatic creation of Host records</li> <li>• <b>In NIOS:</b> existing zone corpxyz.com; no DNS records</li> </ul>	No data for vma	10.10.10.1 vma.corpxyz.com	<b>Zone:</b> corpxyz.com	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)
<ul style="list-style-type: none"> <li>• Add new VM (vma) on Cloud Platform appliance</li> <li>• Automatic creation of Host records</li> <li>• <b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by vDiscovery or admin</i>)</li> </ul>	No data for vma	10.10.10.1 vma.corpxyz.com	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)

Actions and Conditions	Cloud Platform Data before vDiscovery	Cloud Platform Data after vDiscovery	NIOS Data before vDiscovery	NIOS Data after vDiscovery
<ul style="list-style-type: none"> <li>• Add new interface to existing VM (vma) with the same discovered name on Cloud Platform appliance</li> <li>• Automatic creation of Host records</li> <li>• <b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by vDiscovery</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com	10.10.10.1 vma.corpxyz.com 10.10.10.2 vma.corpxyz.com	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1, 10.10.10.2)
<ul style="list-style-type: none"> <li>• Add new interface to existing VM (vma) with the same discovered name on Cloud Platform appliance</li> <li>• Automatic creation of Host records</li> <li>• <b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by admin</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com	10.10.10.1 vma.corpxyz.com 10.10.10.2 vma.corpxyz.com	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)
<ul style="list-style-type: none"> <li>• Add a new interface to existing VM (vma) with a different discovered name (vma-if2) on the Cloud Platform appliance</li> <li>• Automatic creation of Host records</li> <li>• <b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by vDiscovery</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com	10.10.10.1 vma.corpxyz.com 10.10.10.2 vma-if2.corpxyz.com	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1) <b>Host record:</b> vma-if2.corp1.com (10.10.10.2)
<ul style="list-style-type: none"> <li>• Add a new interface to existing VM (vma) with a different discovered name (vma-if2) on the Cloud Platform appliance</li> <li>• Automatic creation of Host records</li> <li>• <b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by admin</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com	10.10.10.1 vma.corpxyz.com 10.10.10.2 vma-if2.corpxyz.com	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1) <b>Host record:</b> vma-if2.corp1.com (10.10.10.2)

Actions and Conditions	Cloud Platform Data before vDiscovery	Cloud Platform Data after vDiscovery	NIOS Data before vDiscovery	NIOS Data after vDiscovery
<ul style="list-style-type: none"> <li>Remove existing VM (vma) on Cloud Platform appliance</li> <li>Automatic creation of Host records</li> <li><b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by vDiscovery</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com	No data for vma	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)	<b>Zone:</b> corpxyz.com
<ul style="list-style-type: none"> <li>Remove existing VM (vma) on Cloud Platform appliance</li> <li>Automatic creation of Host records</li> <li><b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by admin</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com	No data for vma	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)
<ul style="list-style-type: none"> <li>Remove existing interface (10.10.10.2) from VM (vma) with different discovered name (vma-if2) on Cloud Platform appliance</li> <li>Automatic creation of host records</li> <li><b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by vDiscovery</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com 10.10.10.2 vma-if2.corp1.com	10.10.10.1 vma.corpxyz.com	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1) <b>Host record:</b> vma-if2.corp1.com (10.10.10.2)	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)
<ul style="list-style-type: none"> <li>Remove existing interface (10.10.10.2) from VM (vma) with different discovered name (vma-if2) on Cloud Platform appliance</li> <li>Automatic creation of host records</li> <li><b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by admin</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com 10.10.10.2 vma-if2.corp1.com	10.10.10.1 vma.corpxyz.com	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1) <b>Host record:</b> vma-if2.corp1.com (10.10.10.2)	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1) <b>Host record:</b> vma-if2.corp1.com (10.10.10.2)

Actions and Conditions	Cloud Platform Data before vDiscovery	Cloud Platform Data after vDiscovery	NIOS Data before vDiscovery	NIOS Data after vDiscovery
<ul style="list-style-type: none"> <li>Update record name (from vma to vm1) for the existing interface (10.10.10.1) on Cloud Platform appliance</li> <li>Automatic creation of host records</li> <li><b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by vDiscovery</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com	10.10.10.1 vm1.corpxyz.com	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)	<b>Zone:</b> corpxyz.com <b>Host record:</b> vm1.corpxyz.com (10.10.10.1)
<ul style="list-style-type: none"> <li>Update record name (from vma to vm1) for the existing interface (10.10.10.1) on Cloud Platform appliance</li> <li>Automatic creation of host records</li> <li><b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by admin</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com	10.10.10.1 vm1.corpxyz.com	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1) <b>Host record:</b> vm1.corpxyz.com (10.10.10.1)
<ul style="list-style-type: none"> <li>Automatic creation of host records</li> <li>Change FQDN template from \${discover_name} to \${vm_name}.corpxyz.com</li> <li><b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by vDiscovery</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com vm_name: ABC	10.10.10.1 vm1.corpxyz.com vm_name: ABC	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)	<b>Zone:</b> corpxyz.com <b>Host record:</b> ABC.corpxyz.com (10.10.10.1)
<ul style="list-style-type: none"> <li>Automatic creation of Host records</li> <li>Change FQDN template from \${discover_name} to \${vm_name}.corpxyz.com</li> <li><b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by admin</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com vm_name: ABC	10.10.10.1 vm1.corpxyz.com vm_name: ABC	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1) <b>Host record:</b> ABC.corpxyz.com (10.10.10.1)



Actions and Conditions	Cloud Platform Data before vDiscovery	Cloud Platform Data after vDiscovery	NIOS Data before vDiscovery	NIOS Data after vDiscovery
<ul style="list-style-type: none"> <li>Change vDiscovery task configuration from creation of host record to A and PTR records</li> <li><b>In NIOS:</b> existing zone corpxyz.com; existing Host record (<i>originally created by vDiscovery</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com	10.10.10.1 vma.corpxyz.com	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)	<b>Zone:</b> corpxyz.com <b>A record:</b> vma.corpxyz.com (10.10.10.1)
<ul style="list-style-type: none"> <li>Change vDiscovery task configuration from creation of host record to A and PTR records</li> <li><b>In NIOS:</b> existing zone corpxyz.com; existing host record (<i>originally created by admin</i>)</li> </ul>	10.10.10.1 vma.corpxyz.com	10.10.10.1 vma.corpxyz.com	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1)	<b>Zone:</b> corpxyz.com <b>Host record:</b> vma.corpxyz.com (10.10.10.1) <b>A record:</b> vma.corpxyz.com (10.10.10.1)

## Scheduling vDiscovery Jobs

You can enable the appliance to start a vDiscovery immediately after you configure it, schedule it for a later date and time, or configure a recurring discovery based on a recurrence pattern. Note that all scheduled vDiscovery jobs are executed in queue based on the order of the schedule in the **vDiscovery Job Manager**. Therefore, a scheduled vDiscovery might be delayed if there are other jobs being executed before its scheduled start time.

- For a new vDiscovery job: From the **Data Management** tab, select the **IPAM** tab, then select **vDiscovery** -> **New** from the Toolbar; or from the **Cloud** tab, select **vDiscovery** -> **New** from the Toolbar.  
or  
To modify an existing job: From the **Data Management** tab, select the **IPAM** tab and click **vDiscovery** -> **Discovery Manager** from the Toolbar, or from the **Cloud** tab, select **vDiscovery** -> **Discovery Manager** from the Toolbar. In the **vDiscovery Job Manager** editor, click the Action icon next to a selected job and select **Edit** from the menu.
- In step five of the *vDiscovery Job* wizard, or in the **Schedule** tab of the *vDiscovery Job Properties* editor, complete the following:
  - Enable:** To ensure that the scheduled vDiscovery job takes place, select this checkbox. When you upgrade from a previous version of NIOS, you must select this checkbox after the upgrade to ensure that the previously configured discovery tasks are being executed at the scheduled time.  
If you select **Once**, complete the following:
    - Choose a **Start Date** using the date picker.
    - Time Zone:** Select the time zone for the scheduled time from the drop-down list.  
If you select **Hourly**, complete the following:
      - Schedule every hour(s) at:** Enter the number of hours between each update instance. You can enter a value from 1 to 24.
      - Minutes past the hour:** Enter the number of minutes past the hour. For example, enter 5 if you want to schedule the rule update five minutes after the hour.
      - Time Zone:** Select the time zone for the scheduled time from the drop-down list.
    - If you select **Daily**, you can select either **Every day** or **Every Weekday** and then complete the following:
      - Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
      - TimeZone:** Select the time zone for the scheduled time from the drop-down list. If you select **Weekly**, complete the following:
        - Scheduleeveryweekon:** Select any day of the week.

- **Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
  - **TimeZone:** Select the time zone for the scheduled time from the drop-down list. If you select **Monthly**, complete the following:
    - **Schedule the day of the month:** Enter the day of the month and the monthly interval. For example, to schedule the rule update on the first day after every 2 months, you can enter Day 1 every 2 month(s).
    - **Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
    - **TimeZone:** Select the time zone for the scheduled time from the drop-down list.
3. Save the vDiscovery job. The appliance lists all vDiscovery jobs in the **vDiscovery Job Manager**, from which you can manage jobs that have not been executed, including modifying selected jobs or deleting some.

## Managing vDiscovery Jobs

You can view all configured vDiscovery jobs or modify some of the settings for selected ones in the **vDiscovery Job Manager**. You can also add a new vDiscovery job by clicking the Add icon.

To view or modify vDiscovery jobs:

1. From the **Data Management** tab, select the **IPAM** tab, then select **vDiscovery** -> **Discovery Manager** from the Toolbar; or from the **Cloud** tab, select **vDiscovery** -> **Discovery Manager** from the Toolbar.
2. The appliance displays the following information in the **vDiscovery Job Manager**:
  - **Name:** The name of the vDiscovery job.
  - **Status:** The current status of the vDiscovery job. Grid Manager displays an icon and descriptive information about the status. You can hover your mouse over the icon to view the current status, as follows:
    - **Job created:** The configured job has been created.
    - **Job starting:** Starting the configured job.
    - **Job in progress:** The job is being executed at the moment.
    - **Job completed:** The job is completed successfully.
    - **Cancelled:** The job was cancelled while it was being executed.
    - **Error while running job:** The job has failed.  
If the "ERROR: PycURL error: (60, 'SSL certificate problem: unable to get local issuer certificate')" error message is displayed, it means that the certificate has expired or is invalid. You need to remove the expired or invalid certificate and upload a new one. If the error is displayed for an AWS vDiscovery job, download the certificates from <https://www.amazontrust.com/repository/> and upload them. If it is displayed for an Azure vDiscovery job, follow the instructions in the Azure documentation to generate a root certificate for Azure and upload it to NIOS. For information about uploading certificates, see [Managing Certificates](#).
  - **Schedule:** Displays the configured schedule for the selected job.
  - **Public IP's Network View:** The name of the network view to which discovered data for public IP addresses belongs.
  - **Private IP's Network View:** The name of the network view to which discovered data for private IP addresses belongs.
  - **Member:** The Grid member that performs the vDiscovery job.
  - **Last Run:** The timestamp when the selected vDiscovery job was last performed.
  - **Comment:** Comments added when the vDiscovery job was created
3. Click the Action icon

next to a selected job, and you can do one of the following:

- **Edit:** Modify the vDiscovery job settings in the *vDiscovery Job Properties* editor. The editor displays the following tabs: **General**, **Endpoint**, **Network View**, **Data Consolidation**, and **Schedule**. Click the tab that contains information you want to modify, and then modify the information as described in Configuring vDiscovery Jobs above.
- **Delete:** Remove the vDiscovery job from the list.
- **Start:** Start vDiscovery now for the selected job. The appliance might display an AMQP error when you first start a newly created vDiscovery job. This is due to the start-on-demand mechanism experiencing a delay in executing the job. Wait a few seconds and start the job again.
- **Stop:** Stop and terminate the vDiscovery job that is in progress. You cannot resume this discovery once you stop it. All discovered data remains intact in the database.

4. Click **Close** to close the **vDiscovery Job Manager**.

## Viewing Discovered Data

After an IP discovery and a vDiscovery, you can view discovered data in the Discovered Data section when you drill down to a specific IP address in the Net Map, Net List, IP Map or IP List panel. For information about IP Map and IP List, see [Viewing and Managing IPv4 Addresses](#).

You can also view discovered data for Cloud specific VMs in the Cloud tab. Note that you must install the Cloud Network Automation license on the Grid Master to view the Cloud tab.

## Viewing IPAM Discovered Data

To view discovered data in the **IPAM** tab:

1. In the top navigation bar, select the network view to which discovered data belongs.
2. From the **Data Management** tab, select the **IPAM** tab.
3. In the IPAM home page, select the network you want to view discovered data and click the link.
4. In the IP List of IP Map panel, click the IP address link to drill down to the IP address properties panel.
5. In the **Discovered Data** section, click the Configure icon to specify data to be displayed in the **Discovered Data** tab.
6. Select fields from the **Available** list and click the arrow key to move them to the **Selected** list and vice versa. You can use CTRL+Click or SHIFT+Click to select multiple fields at one time.  
You can also change the order in which the fields are displayed by moving the fields up and down in the **Selected** list. To move a field up in the list, select it and click the Up arrow. To move a field down, select it and click the Down arrow.
7. Click the Configure icon again, and the **Discovered Data** section displays the fields you have specified. Note the fields you select appears for all IP addresses in the network.

Depending on the source of the discovered data, when you modify certain DNS and DHCP objects, Grid Manager can display some of the following discovered data (if any) in the **Discovered Data** tab:

- **NetBIOSName**: The name returned in the NetBIOS reply or the name you manually register for the discovered host.
- **OS**: The operating system of the detected host or virtual entity. The OS can be one of the following:
  - **Microsoft** for all discovered hosts that have a non-null value in the MAC addresses using the NetBIOS discovery method.
  - A value that a TCP discovery returns.
  - The OS of a virtual entity on a vSphere server.
- **Discovered MAC Address**: The discovered MAC address for the host. This is the unique identifier of a network device. The discovery acquires the MAC address for hosts that are located on the same network as the Grid member that is running the discovery. This can also be the MAC address of a virtual entity on a specified vSphere server.
- **AP Name**: The name of the access point for the device. This column is displayed only for wireless devices. If discovered by NetMRI, the value will be populated through IPAM Sync as well.
- **AP IP Address**: The IP address of the access point for the device. This column is displayed only for wireless devices. If discovered by NetMRI, the value will be populated through IPAM Sync as well.
- **SSID**: The unique name of the WLAN (Wireless Local Area Network). If discovered by NetMRI, the value will be populated through IPAM Sync as well.
- **Discovered DUID**: For IPv6 address only. The DHCP unique identifier of the discovered host. This is an optional field, and data might not be included.
- **Last Discovered**: The timestamp when the IP address was last discovered.
- **First Discovered**: The timestamp when the IP address was first discovered.
- **Task Name**: The name of the task that collected the discovered data. It is usually the ID or task name that collected the data. It is defined on the corresponding NetMRI appliance when you import the discovered data to the NIOS appliance. The task name should be defined in thevDiscoverytask managerforvDiscovery.
- **Discovered Name**: The name of the network device associated with the discovered IP address.
- **Discoverer**: Specifies whether the IP address was discovered by NetMRI or NIOS discovery process.

If you imported data from NetMRI appliances, Grid Manager displays the following information, if available. For information about the data imported from NetMRI appliances, see [Integrating Discovered Data From NetMRI](#).

- **Attached Device Description:** A textual description of the switch that is connected to the end device.
- **Attached Device Address:** The IPv4 or IPv6 address of the switch that is connected to the end device.
- **Attached Device Model:** If a reverse lookup was successful for the IP address associated with this switch, the device model is displayed here.
- **Attached Device Name:** If a reverse lookup was successful for the IP address associated with this switch, the host name is displayed here.
- **Attached Device Port Description:** A textual description of the switch port that is connected to the end device.
- **Attached Device Port Name:** The name of the switch port or port channel connected to the device. For Cisco devices with virtual port channel configured, this field also displays the list of physical interfaces that form the virtual port channel. For more information, see [IP List Neighbor Information](#).
- **Attached Device Vendor:** The vendor name of the switch port connected to the end device.
- **Attached Device Port:** The number of the switch port connected to the end device.
- **Attached Device Type:** Identifies the switch that is connected to the end device.
- **Attached Device Location:** The physical location of the network device to which the end host is connected/ attached, as detected from the device during discovery.
- **Attached Device Contact:** The contact details for the network device to which the End Host is connected/ attached, as detected from the device during discovery.
- **Device Vendor:** The vendor name of the end device.
- **Device Model:** The device model of the end device.
- **Device Location:** The physical location of the network device on which the IP Address is configured, as detected from the device during discovery.
- **Device Contact:** The contact details for the network device on which the IP Address is configured, as detected from the device during discovery.
- **Device Management IP:** The IPv4 or IPv6 address of the end device that is connected to the switch.
- **Device Port Type:** The port type for the end device.
- **Device Port Name:** The port name for the end device.
- **Device Type(s):** Identifies the device type.
- **Port Duplex:** The negotiated or operational duplex setting of the switch port connected to the end device. You can modify this in the IPv6 fixed address and AAAA record editors.
- **Port Link:** The link status of the switch port connected to the end device. Indicates whether it is connected.
- **Port Speed:** The interface speed, in Mbps, of the switch port. You can modify this in the IPv6 fixed address and AAAA record editors.
- **Port Type:** The switch port type.
- **Port Status:** The operational status of the switch port. Indicates whether the port is up or down.
- **Open Port(s):** Ports that are open.
- **VLAN Name:** The name of the VLAN of the switch port.
- **VLAN ID:** The ID of the VLAN of the switch port.

For IP addresses discovered through a vDiscovery, Grid Manager displays the following additional information, if available:

- **Virtual Host Adapter:** The name of the physical network adapter through which the virtual entity is connected to the appliance.
- **Virtual Datacenter:** The name of the vSphere datacenter or container to which the virtual entity belongs.
- **Virtual Cluster:** The name of the VMware cluster to which the virtual entity belongs.
- **Virtual Entity Name:** The name of the virtual entity.
- **Virtual Entity Type:** The virtual entity type. This can be blank or one of the following: Virtual Machine, Virtual Host, or Virtual Center. Virtual Center represents a VMware vCenter server.
- **Virtual Host:** The name of the VMware server on which the virtual entity was discovered.
- **Virtual Switch:** The name of the switch to which the virtual entity is connected.
- **Virtual Machine Name:** The name of the VM instance.
- **Virtual Machine ID:** The ID of the VM.
- **Virtual Machine Tenant ID:** The tenant ID to which the VM belongs.
- **Virtual Machine Port Group:** The port group to which the VM belongs.
- **Attached Virtual Switch Name:** The name of the virtual switch to which the VM is connected.
- **Attached Virtual Switch ID:** The ID of the virtual switch to which the VM is connected.
- **Attached Virtual Switch Type:** The type of the virtual switch. This can be standard or distributed.

- **Attached Virtual Switch IPv6 Enabled:** Indicates that virtual switch that has IPv6 enabled.
- **Attached Virtual Port Name:** The name of the virtual adapter on the virtual switch to which the VM is connected.
- **Attached Virtual Port MAC Address:** The MAC address of the virtual adapter on the virtual switch to which the VM is connected.
- **Attached Virtual Port Link Status:** The link status of the virtual adapter on the virtual switch to which the VM is connected.
- **Attached Virtual Port Configured Speed:** The configured port speed of the virtual adapter on the virtual switch to which the VM is connected.
- **Attached Virtual Port Configured Mode:** The configured mode of the virtual adapter on the virtual switch to which the VM is connected.
- **Attached Virtual Port Configured Speed:** The actual port speed of the virtual adapter on the virtual switch to which the VM is connected.
- **Attached Virtual Port Configured Mode:** The actual mode of the virtual adapter on the virtual switch to which the VM is connected.
- **Network Segment Type:** The type of network segment to which the VM is connected.
- **Network Segment Name:** The name of the network segment to which the VM is connected.
- **Network Segment ID:** The ID of the network segment to which the VM is connected.
- **Network Segment Port Group:** The port group of the network segment to which the VM is connected.
- **Network Segment Available Ports:** The number of available ports reported by the virtual switch to which the VM is connected.
- **Attached Virtual Switch VTEP Type:** The type of VTEP (Virtual Tunnel Endpoint) in the virtual switch that is connected to the VM. This can be VXLAN or STT.
- **Attached Virtual Switch VTEP IP:** The IP address of the VTEP in the virtual switch that is connected to the VM.
- **Attached Virtual Switch VTEP Port Group:** The port group of the VTEP in the virtual switch that is connected to the VM.
- **Attached Virtual Switch VTEP VLAN:** The VLAN of the VTEP in the virtual switch that is connected to the VM.
- **Attached Virtual Switch VTEP DHCP Server:** The DHCP server of the VTEP in the virtual switch that is connected to the VM.
- **Attached Virtual Switch VTEP Multicast:** The multicast address of the VTEP in the virtual switch that is connected to the VM.
- **Physical Host IP Address:** The IP address of the physical node on which the VM is hosted.
- **Physical Host Name:** The name of the physical node on which the VM is hosted.
- **Physical Host MAC Address:** The MAC address of the physical node on which the VM is hosted.
- **Physical Host CIDR Subnet:** The subnet mask of the physical node on which the VM is hosted.
- **Physical Host 's NIC Names:** The list of all physical port names used by the virtual switch on the physical node on which the virtual machine is hosted. Valid values are eth1, eth2, eth3, and so on.

## Viewing Discovered Data for Cloud VMs

You can view discovered data for cloud specific VMs after performing a vDiscovery job. Note that you must have at least read-only permission to "All Tenants" and "All Network Views" to view the **VMs** tab.

To view discovered data for cloud VMs, see [Viewing All Cloud VMs](#).

## Managing Discovered Data

In addition to viewing the discovered data, you can perform the following to manage the discovered data:

- Manage an unmanaged address by adding it to a host, converting it to managed data, or clearing its unmanaged status. Note that for Cloud Network Automation, you can convert discovered data to NIOS objects within a delegated scope.
- Resolve conflicting addresses.
- Clear discovered data for a network view, network, IP address, or cloud tenant.
- You can also clear all discovered data collected by a specified vDiscovery job.

This section covers the following:

- [Managing Unmanaged Data](#)
  - [Adding to an Existing Host](#)

- [Converting Unmanaged Data](#)
- [Clearing Unmanaged Data](#)
- [Resolving Conflicting Addresses](#)
  - [Resolving DHCP Lease Conflicts](#)
  - [Resolving Fixed Address Conflicts](#)
  - [Resolving DHCP Range Conflicts](#)
  - [Resolving Host Conflicts](#)
  - [Resolving VM Affiliation Conflicts](#)
  - [Resolving MAC Address Conflicts](#)
- [Clearing Discovered Data](#)
- [Clearing All Discovery Data](#)

## Managing Unmanaged Data

You can manage unused and unmanaged addresses by performing one of the following:

- Add to an existing host, as described in [Adding to an Existing Host](#) below.
- Convert to a fixed address, host, A record, or PTR record, as described in [Converting Unmanaged Data](#) below. Note that for Cloud Network Automation, you can convert discovered data to NIOS objects within a delegated scope.
- Clear the unmanaged status, as described in [Clearing Unmanaged Data](#) below.



### Note

You cannot convert unmanaged IP addresses that are being served by Microsoft DHCP servers to host records.

## Adding to an Existing Host

You can add an unmanaged address, including all its information, to an existing host. You can select the desired host to which you want to add the unmanaged address.

To add an unmanaged address to an existing host, perform the following:

1. From the *IP Map* or *List* panel, select an unmanaged address you want to add to a host, and then click **Add** -> **Add to Existing Host** from the Toolbar.
2. In the *Select Host* dialog box, select a host from the table. You can also search for a host using filters or the Go to function. For information Using filters and Using the go to Function , see [Finding and Restoring Data](#). Click the Select icon to select the desired host.  
Note that depending on the page size configuration, the search results are limited to the page size that you set. If the search results exceed the page size limit, the appliance displays an error message to inform you to refine your search criteria or to change the page size limit. In the *Host Record* editor, complete the information.
3. Save the configuration, and then click **Restart** if it appears at the top of the screen.

## Converting Unmanaged Data

You can convert an unmanaged address to a host, an A or AAAA record, a PTR record, or a fixed address. To convert an unmanaged address:

1. In the *IP Map* or *List* panel, select an unmanaged address you want to convert, and then select **Convert** from the Toolbar.
2. In the drop-down list, select the type of address to which you want to convert the unmanaged address. For IPv4 addresses, you can select **To Host**, **To A Record**, **To PTR Record**, or **To Fixed Address**. For IPv6 addresses, you can select **To Host**, **To AAAA**, **To PTR Record**, or **To IPv6 Fixed Address**.  
Depending on the record type you select, Grid Manager displays the corresponding editor. It also populates the attributes of the unmanaged address in the editor. Enter the appropriate information in the editor.
3. Save the configuration, and then click **Restart** if it appears at the top of the screen.





#### Note

After the conversion, the status of the unmanaged address changes to **Used**.

The following are some conditions for a conversion:

- **A and AAAA records:** You must select a DNS zone when converting an unmanaged address to an A or AAAA record.
- **PTR record:** You must select a DNS zone when converting an unmanaged address to a PTR record.
- **IPv4 and IPv6 Fixed Address:** Grid Manager displays a confirmation dialog box to ensure that you want to create a fixed address for the unmanaged address.
- **IPv4 and IPv6 Host record:** You can use the unmanaged address to enable a host record for DNS or DHCP.

## Clearing Unmanaged Data

Unmanaged objects are objects that are not configured for DNS or DHCP and do not have any corresponding NIOS objects such as fixed addresses, DNS records, or host records. You can clear unmanaged data if you do not want it to appear in the discovered data. Any unmanaged data that was discovered in a previous discovery can be removed. When you clear an unmanaged IP address, the status of the IP address changes to **Unused**.

Note that for cloud deployments, you can remove an unmanaged VM only when all the VM related properties are removed first. When the VM is removed, the tenant pointing to this VM will also be removed automatically. For information about cloud deployments, see [Deploying Cloud Network Automation](#).

To clear unmanaged data, complete the following steps:

### Method 1

1. **For IP addresses:** From the **Data Management** tab -> **IPAM** tab, select an unmanaged IP address in the *IPMap* or *List* panel.  
**For networks:** From the **Data Management** tab -> **IPAM** tab -> *IPMap* or *List* panel, select a network in which you want to clear all unmanaged addresses.  
**For Cloud tenants:** From the **Cloud** tab -> **Tenants** tab, select a tenant for which you want to clear unmanaged data.  
**For Cloud VMs:** From the **Cloud** tab -> **VMs** tab, select a VM for which you want to clear unmanaged data.
2. Click **Clear** -> **Clear Unmanaged Data** or **Clear All Unmanaged Data** from the Toolbar.
3. In the *ClearUnmanagedData* confirmation dialog box, click **Yes**. The appliance clears data that has no corresponding NIOS objects such as fixed addresses, DNS records, or host records.

### Method 2

1. From the **Cloud** tab -> **VM** tab, select **Discovery Manager** from the Toolbar.
2. In the *vDiscoveryJobManager* dialog, click the Action icon  
next to the selected vDiscovery job, and then select **Clear Unmanaged Data**.
3. In the *ClearUnmanagedData* confirmation dialog box, click **Yes**. The appliance clears data that has no corresponding NIOS objects such as fixed addresses, DNS records, or host records.



#### Note

After you clear all the unmanaged data, you should navigate to the **Data Management** tab and clear all the associated networks. When you clear unmanaged addresses in a given network view, all unmanaged IPv4 and IPv6 addresses of all networks in the network view are cleared. When you select an entire network or a specific network in the *IP Map* or *List* panel, all the unmanaged addresses in the network are cleared. After you clear the unmanaged data, the status of the IP addresses changes to **Unused**.

## Resolving Conflicting Addresses

Conflicts happen when discovered data does not match existing IP address data. The *IP Map* panel and the **Cloud** tab -> **VMs** tab display conflicting addresses with a red background. The *IP List* panel displays **Conflict** as the status for all

conflicting addresses. For objects that have multiple conflicts, Grid Manager lists each of them in a bulleted list in the *Resolve multiple conflicts* dialog. You can select a conflict, in any order, to begin resolving each issue. After you resolve the selected issue, Grid Manager returns to the *Resolve multiple conflicts* dialog so you can resolve other issues. Depending on the conflict, you can do one of the following to resolve it:

- For a DHCP lease conflict, you can clear the existing lease and create either a fixed address or a reservation for the IP address. You can also keep the existing data and clear the discovered data. For more information, see *Resolving DHCP Lease Conflicts* below.
- For a fixed address conflict, you can either keep the existing fixed address data or update the existing data with the discovered data. For more information, see *Resolving Fixed Address Conflicts* below.
- For a DHCP range conflict, you can create a fixed address, create a reservation, or clear the discovered data. For more information, see *Resolving DHCP Range Conflicts* below.
- For a host conflict, you can either keep the existing host record data or update the existing data with the discovered data. For more information, see *Resolving Host Conflicts* below.
- For a VM affiliation conflict, you can either update all the displayed objects to be affiliate with the discovered data or keep the current affiliation and clear the conflict. For more information, see *Resolving VM Affiliation Conflicts* below.

You must resolve conflicting addresses individually. You cannot resolve multiple conflicts at the same time.



#### Note

After the conflict is resolved, the status of the IP address changes depending on how you resolved the conflict.

To resolve a conflict:

1. In the *IP Map* or *List* panel, select a conflicting address, and then click **Resolve Conflict** from the Toolbar.
2. The *Resolve Conflict* dialog box displays the reason of the conflict and lists the existing information and discovered information of the address in the **Description** field. Depending on the type of conflict, the appliance displays the corresponding resolution options. You can compare the existing and discovered data and decide how you want to resolve the conflict.

### Resolving DHCP Lease Conflicts

When an IP address has a DHCP lease and the discovered MAC address is in conflict with the existing MAC address, the IP address has a DHCP lease conflict.

To resolve a DHCP lease conflict, perform the following:

1. In the *Resolve Conflict* or *Resolve multiple conflicts* dialog, Grid Manager displays the nature of the conflict and the discovered data versus the current data. Select one of the following to resolve the conflict:
  - **Clear lease and create fixed address from discovered data:** Clears the existing DHCP lease and creates a fixed address with the discovered data. The *Fixed Address* editor appears with the discovered data populated.
  - **Clear lease and create a reservation from discovered data:** Clears the existing DHCP lease and creates a new reservation using the discovered data. The *Reservation* editor appears with the discovered data populated. This option does not apply to leases served by Microsoft DHCP servers because they do not support Infoblox reservations.
  - **Keep the existing and ignore this conflict:** Keeps the current DHCP lease for the address and ignores the lease conflict.
2. Click **OK** or **Resolve** (when you have multiple conflicts). If you have multiple conflicts, Grid Manager returns to the *Resolve multiple conflicts* dialog so that you can resolve other conflicts.

### Resolving Fixed Address Conflicts

When the discovered MAC address of an IPv4 address does not match with its existing MAC address, or when the DUID of an IPv6 address does not match with its existing DUID, the IP address has a fixed address conflict.

To resolve a fixed address conflict, perform the following:



1. In the *Resolve Conflict* or *Resolve multiple conflicts* dialog, Grid Manager displays the nature of the conflict and the discovered data versus the current data. Select one of the following to resolve the conflict:
  - **Keep fixed address and clear discovered data:** Keeps the existing fixed address and clears the discovered data.
  - **Update fixed address with discovered data:** Updates the existing fixed address data with the discovered data.
2. Click **OK** or **Resolve** (when you have multiple conflicts). If you have multiple conflicts, Grid Manager returns to the *Resolve multiple conflicts* dialog so that you can resolve other conflicts.

### Resolving DHCP Range Conflicts

When an IP address is in a DHCP range and does not match an existing DHCP lease, fixed address, or exclusion range and it shows an active state during a discovery, the IP address has a DHCP range conflict.

To resolve a DHCP range conflict:

1. In the *Resolve Conflict* or *Resolve multiple conflicts* dialog, Grid Manager displays the nature of the conflict and the discovered data versus the current data. Select one of the following to resolve the conflict:
  - **Create a fixed address:** Creates a fixed address with the discovered data. If the fixed address is served by a Microsoft server, but is outside of a scope, you must then navigate to the *Fixed Address* editor and assign the fixed address to the appropriate Microsoft server.
  - **Create a reservation:** Creates a reservation with the discovered data. This creates an Infoblox reservation, and therefore cannot be used for IP addresses that are served by Microsoft servers. Note that you cannot convert an IPv6 address to a reservation.
  - **Clear discovered data:** Clears the discovered data and no object is created for the IP address.
2. Click **OK** or **Resolve** (when you have multiple conflicts). If you have multiple conflicts, Grid Manager returns to the *Resolve multiple conflicts* dialog so that you can resolve other conflicts.

### Resolving Host Conflicts

When the MAC address of an IPv4 address that belongs to a host record does not match with its existing MAC address, or when the DUID of an IPv6 address that belongs to a host record does not match with its existing DUID, the IP address has a host conflict.

1. In the *Resolve Conflict* or *Resolve multiple conflicts* dialog, Grid Manager displays the nature of the conflict and the discovered data versus the current data. Select one of the following to resolve the conflict:
  - **Keep host record and clear discovered data:** Keeps the existing data and clears the discovered data.
  - **Update host record with discovered data:** Updates the existing host record data with the discovered data.
2. Click **OK** or **Resolve** (when you have multiple conflicts). If you have multiple conflicts, Grid Manager returns to the *Resolve multiple conflicts* dialog so that you can resolve other conflicts.

### Resolving VM Affiliation Conflicts

When an IP address contains objects that are affiliated with a VM that is not the same as the discovered VM, this IP address has a VM affiliation conflict.

To resolve a VM affiliation conflict:

1. In the *Resolve Conflict* or *Resolve multiple conflicts* dialog, Grid Manager displays the nature of the conflict and the discovered affiliation versus the current affiliation. Select one of the following to resolve the conflict:
  - **Update all the displayed objects to be affiliate to the discovered affiliation:** Updates all associated objects in this IP address to affiliate with the discovered affiliation.
  - **Keep the current affiliation(s) and clear the conflict:** Keeps the existing data and clears the conflict.
2. Click **OK** or **Resolve** (when you have multiple conflicts). If you have multiple conflicts, Grid Manager returns to the *Resolve multiple conflicts* dialog so that you can resolve other conflicts.

### Resolving MAC Address Conflicts

When the MAC address of an existing IP address does not match the MAC address of the discovered data, the object has a MAC address conflict.

To resolve a MAC address conflict:

1. In the *Resolve Conflict* or *Resolve multiple conflicts* dialog, Grid Manager displays the nature of the conflict and the discovered data versus the current data. Select one of the following to resolve the conflict:
  - **Change the configured MAC address to be the same as the discovered MAC address:** Changes the MAC address to the discovered MAC address.
  - **Keep fixed address and ignore this conflict:** Keeps the fixed address and ignores the discovered data.
2. Click **OK** or **Resolve** (when you have multiple conflicts). If you have multiple conflicts, Grid Manager returns to the *Resolve multiple conflicts* dialog so that you can resolve other conflicts.

## Clearing Discovered Data

You can clear previously discovered managed data for selected IPv4 or IPv6 networks. This action is useful, for example, if your network topology has changed since the last discovery and you want to discover new data on the network or cloud platform. You may perform this action whether or not the network is in a managed or unmanaged state.



### Note

This action clears only the discovered data that is supported in the **Discovered Data** section for an IP address. It does not clear any NIOS objects or information such as tenants, networks, or VMs for a cloud platform. If a discovered IP address has the same IP as an existing NIOS object (such as a fixed address, DNS record, or host record), the appliance removes this IP address.

To clear discovered data:

1. In the *IP Map* or *List* panel, select a network, and then click **Clear** -> **Clear Discovered Data** from the Toolbar.
2. In the *Clear Discovered Data* dialog box, click **Yes**.
3. Navigate to the **Data Management** tab → **DNS** tab, delete all the associated zones.
4. After you clear all the discovered data, you should navigate to the **Data Management** tab and clear all the associated networks.

You can also clear discovered data on all networks in a network view by performing the following steps:

1. In the *IP Map* or *List* panel, select a network, and then click **Clear** -> **Clear Discovered Data** from the Toolbar.
2. In the *Clear Discovered Data* dialog box, click **Yes**.
3. Navigate to the **Data Management** tab → **DNS** tab, delete all the associated zones.
4. After you clear all the discovered data, you should navigate to the **Data Management** tab and clear all the associated networks.



### Note

When you clear all discovered data in a given network view, all imported discovered data for managed addresses, in all IPv4 and IPv6 networks in the network view, are cleared.

You can also clear discovered data for a specific discovery job, as follows:

1. From the **Cloud** tab -> **VM** tab, select **DiscoveryManager** from the Toolbar.
2. In the *vDiscoveryJobManager* dialog, click the Action icon  
next to the selected vDiscovery job, and then select **ClearDiscoveredData**.
3. In the *ClearDiscoveredData* dialog box, click **Yes**. The appliance clears all the discovered managed data that is collected by the specified vDiscovery job.
4. Navigate to the **Data Management** tab → **DNS** tab, and then delete all the associated zones.
5. After you clear all the discovered data, you should navigate to the **Data Management** tab and clear all the associated networks.



#### Note

If you delete an associated zone before you clear the discovered data, the **Clear** option gets disabled from the **Cloud** tab -> **VM** tab. The **Clear** option can be accessed from the **Cloud** tab -> **Tenants** tab. This is applicable when you use WAPI calls to clear all discovery data.

## Clearing All Discovery Data

You can clear all the discovered data, whether managed or unmanaged, for a specific vDiscovery job. This action removes only the discovered data, not the associated NIOS objects, collected by the specified vDiscovery job only. It does not remove any discovered data collected through Network Insight or other non vDiscovery tasks.



#### Note

If the same data is collected by the specified vDiscovery job and another non vDiscovery job such as Network Insight or IP discovery, the discovered data remains intact and will not be removed.

To clear all discovered data for a specific vDiscovery job, perform the following:

1. From the **Cloud** tab -> **VM** tab, select **Discovery Manager** from the Toolbar.
2. In the *vDiscoveryJobManager* dialog, click the Action icon  
next to the selected vDiscovery job, and then select **Clear All Discovery Data**.
3. In the *ClearDiscoveredData* dialog box, click **Yes**. The appliance clears all the discovered managed and unmanaged data that is discovered by the specified vDiscovery job.



#### Note

If you delete all the networks from the **Data Management** tab before you clear all the discovered data, the data gets cleared only from the NIOS user interface and not from the NIOS database which results in stale VM to be fetched when WAPI calls are made.

## Integrating Discovered Data from NetMRI

The NetMRI appliances discover and track IPv4 and IPv6 network devices and provide information about the discovered IP addresses. You can integrate IPv4 and IPv6 discovered data into the NIOS appliance database, and then view the data in the IP List panel of Grid Manager as well as in the **Discovered Data** tab of certain IPAM object editor. When you start synchronization of discovered data from the NetMRI appliance, only the available IP addresses in the network are discovered. The imported data from the NetMRI appliance can only be modified or deleted by the specific synchronization. When you import the discovered data again for the same network and if some of the information in the network is changed, the newly discovered data and the originally discovered data (now old data) co-exist in the database. When you import the data again, the old data is deleted or modified, or the new data is added depending on the information discovered.

For information about NetMRI IP discovery and how to import discovered data from a NetMRI appliance to the NIOS appliance, refer to [Executing NIOS IPAM Sync](#) in the Infoblox NetMRI documentation.



#### Note

- NIOS does not import IPv6 leases that contain prefixes and link-local IPv6 addresses. This data is discarded during an import.
- After an IPAM synchronization, CSV import errors if any are logged in a separate file named `discovery_csv_error.log.xxxxxx` located at `/infoblox/var/discovery_csv_error`

The appliance can import the following IPv4 and IPv6 data that NetMRI discovers:

- **IP Address:** The discovered IPv4 or IPv6 address.
- **Discovered MAC Address:** The MAC address of the discovered host.
- **Last Discovered:** The date and time the IP address was last discovered.
- **NetBIOS Name:** The name returned in the NetBIOS reply or the name that you manually register for the discovered host.
- **OS:** The operating system of the detected host.
- **First Discovered:** The date and time the IP address was first discovered.
- **Discoverer:** Specifies whether the IP address was discovered by a NetMRI discovery process.
- **Discovered Name:** The name of the network device associated with the discovered IP address.
- **Attached Device Description:** A textual description of the switch that is connected to the end device.
- **Attached Device Address:** The IP address of the switch that is connected to the end device.
- **Attached Device Name:** If a reverse lookup was successful for the IP address associated with this switch, the host name is displayed here.
- **Attached Device Port Description:** A textual description of the switch port that is connected to the end device.
- **Attached Device Port Name:** The name of the switch port or port channel connected to the device. For Cisco devices with virtual port channel configured, this field also displays the list of physical interfaces that form the virtual port channel. For more information, see [IP List Neighbor Information](#).
- **Attached Device Port:** The number of the switch port connected to the end device.
- **Attached Device:** Identifies the switch that is connected to the end device.
- **Port Duplex:** The negotiated or operational duplex setting of the switch port connected to the end device.
- **Port Link:** The link status of the switch port connected to the end device. Indicates whether it is connected.
- **Port Speed:** The interface speed, in Mbps, of the switch port.
- **Port Status:** The operational status of the switch port. Indicates whether the port is up or down.
- **VLAN Name:** The name of the VLAN of the switch port.
- **VLAN:** The ID of the VLAN of the switch port.

## Infoblox Network Insight

This section provides information about Infoblox Network Insight, which unites network discovery for geographically dispersed networks, infrastructure devices, and network assets with the Infoblox IPAM (IP Address Management) solution. Through discovery, Network Insight provides automated, comprehensive, and accurate data about your network devices and their attributes so you can have better visibility in your network infrastructure, including the virtual network infrastructure, and manage it more efficiently. This section also describes how you can use Network Insight to detect and manage information about network infrastructure devices, how to provision and de-provision networks, and how to manage and provision device ports, including switched Ethernet. It includes the following topics:

- [Administrative Permissions for Discovery](#)
- [About Automatic Conversion Rules](#)
- [Accessing Detailed Device Information](#)
- [Adding Discovery Device Support](#)
- [Excluding IP Addresses from Discovery](#)
- [Disabling Discovery for a Network](#)
- [Discovering Devices and Networks](#)
- [Starting Discovery](#)
- [Network Insight Architecture](#)
- [Mapping Discovery Interfaces to Network Views](#)

- [Managing Discovery](#)
- [Conflict Resolution in Network Insight](#)
- [Executing Discovery Diagnostics](#)
- [Discovering VRF Virtual Networks](#)
- [Configuring Discovery Properties](#)
- [Configuring Discovery for SDN and SD-WAN](#)
- [Consolidator and Probes](#)
- [Defining the Discovery Member Type](#)
- [Viewing Discovered Devices and their Properties](#)
- [Port Control Features in Network Insight](#)
- [Defining Blackout Periods](#)
- [Defining Port Configuration Blackout Periods](#)
- [Creating Port Reservations for IPAM Objects](#)
- [Editing Interfaces in a Device](#)
- [Monitoring Device Lifecycle and Vulnerabilities Using Advisor](#)
- [Viewing Discovery Status](#)
- [Provisioning and De-Provisioning Networks](#)
- [Supported Discovery Methods](#)
- [Starting and Stopping the Discovery Service](#)
- [Managing Discovered Data](#)
- [Configuring Automatic VRF Mapping](#)
- [About Network Insight](#)

## Administrative Permissions for Discovery

You can start a discovery and manage discovered data based on your administrative permissions. For more information, see [About Administrative Permissions](#).

Initiating and controlling a discovery requires specific administrative permissions. The following are permission guidelines for initiating and controlling a discovery:

- The **IPAM Discovery Admin** role provides a pre-configured list of permissions by which assigned admin accounts may perform discovery tasks. Administrators with these permissions can initiate and control discovery on any existing network. By default, the **IPAM Discovery Admin** role supports the following permissions:
  - All permissions associated with the Network Discovery feature set (active if you do not have a Network Insight license)
  - Read-Only on all Network Views, network containers, networks and ranges
  - Read-Only on all hosts
  - Read-Only on all Members
  - Read-Write Network Discovery permissions
  - Editing network, network container or range discovery properties: Read-Only for each type. For member assignment, the user also needs additional read-only permission for the assigned member
  - Editing fixed address, host or reservation discovery properties: Read-only for each type
  - Excluding an IP address or an IP Range (from the Network Editor 's **Discovery Exclusions** tab or from the IPAM IP **List** view): read-only permission for the network
  - Defining discovery and port configuration blackouts at the Grid and Network levels
  - **Discover Now** for Network, DHCP Range, IP address or device: Read-Only permission for each.
- If the user does not possess the **Network Discovery** permission, all Network Insight permissions and operations are disabled.
- A **Port Control** permission under the **IPAM Permissions** type allows you to add a Read/Write or Deny permission for a device; a network; a network container; or network view. You can also add a global RW/Deny port control permission for all network views.
  - If the user does not possess the **Port Control** permission, the user cannot provision networks, de-provision networks, perform port reservations or configure interfaces. *All non-Superuser accounts must have the Port Control permission to create port control tasks for any affected objects.*
- Superusers can initiate and control discovery on all networks. Some discovery functions require superuser permissions:

- Grid Discovery properties
- Uploading, Viewing and deletion of device support bundles
- Launching Discovery Diagnostics
- Launching Discovery Status

Similar to Network Discovery, devices and end hosts found through discovery can undergo conversion from unmanaged status to managed status. This entails converting an unmanaged IP address to a host, an A record or AAAA record, a PTR record, or to a fixed address.

- Administrators with read/write permission to a DNS zone or specific record type can convert unmanaged data to a host, fixed address, reservation, A record, or PTR record.
- IPAM Discovery admins can convert unmanaged networks to managed networks and can change discovery settings for networks.
- For unmanaged networks: users may **Delete**, **Convert** (to managed), **Clear Unmanaged Data** and **Clear Discovered Data** if one of the following is true:
  - User has read-write permission for the network
  - or–
  - User has Network Discovery permission plus Read-Only for the network.

After a discovery is complete, the following permission guidelines apply to viewing and managing discovery data:

- Superusers can view and manage all discovered data.
- Administrators with read permission to networks can view all discovery data without editing.
- If a user has read-only permissions for a device's management IP address, the device is visible in the **Data Management** → **Devices** tab.

For more information on configuring admin user accounts and working with permissions, see the sections under [Managing Administrators](#). For information about discovery permissions, see [Administrative Permissions for Network Insight Tasks](#).

## About Automatic Conversion Rules

To automate the conversion of IP addresses of discovered entities from "unmanaged" to "managed" in a specific network view, you can configure conversion rules that Network Insight uses to automatically create new DNS records or update existing data for the discovered IP addresses. Network Insight automatically converts newly discovered IP addresses to host records, A and PTR records, or fixed addresses based on your configuration. You can define templates that Network Insight uses to create new records by using supported variables and functions. For information about supported variables, see [Supported Variables for Templates](#) below.

Note that corresponding DNS zones in a selected network view must already exist in order for Network Insight to add DNS records during the conversion. Otherwise, Network Insight does not add any DNS records and it logs a message to the syslog.

Network Insight automatically adds DNS records based on the following conditions:

- The corresponding DNS zones must already exist in the NIOS database. Network Insight does not automatically create DNS zones for the records.
- To create a PTR record, the corresponding reverse-mapping zone must exist.
- A DNS zone cannot be associated with more than one DNS view. Network Insight does not create DNS records for zones that are associated with multiple DNS views.
- NIOS adds new DNS records only if the discovered\_name for the discovered IP address is available and there is no conflict with information about the associated network view.

On subsequent discovery jobs, if an IP for a VM is removed, the corresponding DNS records are removed accordingly. If the IP for a VM is changed, the IP address in the corresponding DNS record is changed accordingly. If the DNS record name template is changed, all the DNS records are replaced with the DNS records using the new template. All administrative actions for these change are recorded in the audit log. Summary of the changes are logged in the syslog.

### Note

Network Insight updates only records that are created by the Network Insight process. It does not create or update DNS records that are originally created by other admin users.

## Configuring Automatic Conversion Rules

To add automatic conversion rules:

1. From the **Grid** tab, select the **Grid Manager** tab, and click **Edit → Grid Discovery Properties** from the Toolbar.
2. Click the **Conversion Policy** tab and complete the following:
  - **Enable the automatic conversion rules defined for newly discovered IP addresses:** Select this checkbox to enable the automatic conversion of unmanaged IP addresses of newly discovered entities to managed objects in a specific network view. This is disabled by default.
  - **Update discovered data for managed objects:** Select this checkbox if you want the appliance to update discovered data for all corresponding NIOS objects (if they exist in NIOS), such as A records, PTR records, host records, and fixed addresses, with the discovered data. If you do not select this checkbox, the appliance updates only the discovered data for unmanaged objects. None of the managed data will be updated. This checkbox is selected by default, but is disabled if you do not enable the automatic conversion feature.

Click the Add (+) icon and Grid Manager adds a row to the table (this table is enabled only when you enable the automatic conversion feature). Complete the following:

- **Network View:** From the drop-down list, select the network view in which your conversion rule will take effect. Note that this rule applies only to objects in the selected network view. If you have multiple network views, you must configure a separate policy for each network view.
- **Template:** Define a naming template that Network Insight uses to automatically create DNS records for the unmanaged IP addresses in the network view. You can use the following syntax: `${substitution}`, where **substitution** can be a supported variable or function. Note that each IPv6 address substitution is unwrapped into dotted presentation. For information about supported variables and functions, see Supported Variables for Templates below.

For example, when you enter `${discovered_name}.corpxyz.com`

and the `discovered_name` for the asset is

`XYZ`, the DNS name for this IP becomes

`XYZ.corpxyz.com`

. When you enter

`$dev-{ip_address_octet3}.corpxyz.com`

and the IP for the asset

`is2dba::db8::1`

, the DNS name for this IP becomes

`dev-3.corpxyz.com`

. When you enter

`${ip_address[7]}.corpxyz.com`

for an IPv6 address and if the IP for the asset is

`2001:db8:acad::1`

, the DNS name becomes

`b.corpxyz.com`.

You can also use the following functions in the naming template: dashed, reversed, and underscored. For example, when you enter

`${dashed(${ip_address})}-corpxyz.com`

and the IP is 1.2.3.4, the DNS name becomes

`1-2-3-4-corpxyz.com`. When you enter

`${reversed(${ip_address})}-corpxyz.com`

and the IP is

`1.2.3.4`

, the DNS name becomes

`4.3.2.1-corpxyz.com`.

- **Conditions:** Enter the matching conditions for the conversion rule. You can use magic variables, supported variables, operators, and functions in the condition. When Network Insight finds IP addresses that match this condition, it will convert the IP addresses into DNS records (Hosts, A/PTR records, or fixed addresses) based on your selected conversion type. For information about supported parameters, see Supported Conversion Parameters below.

For example, if you want to match IP addresses that do not have an FQDN in the

`discovered_name`

, you can enter this condition:

`${is_fqdn(${discovered_name})} == false AND $`

`{discovered_name} == 'unknown'`. That is, you can use `unknown` as the filter value for IP addresses that do not have FQDN.

If you want to match devices from the network `192.168.1.0/24`

with the name starting with "

`Serial0`

", you can enter this condition:

`${ip_belongs_to("192.168.1.0/24")} == true AND`

`${discovered_name} like "Serial0"`.

**ConversionType:** From the drop-down list, select the DNS record type that you want Network insight to convert the unmanaged IP addresses into. You can convert an unmanaged IP into **Host**, **A/PTR**, or **FixedAddress**. When you select **A/PTR**, Network Insight converts each IP into A and PTR records simultaneously.

**Comment:** Enter description about this policy to distinguish it from others. For example, if the policy is used to identify and convert IP addresses with

`discovered_name`

that does not contain an FQDN, you can enter "No FQDN in discovered\_name." as the comment to remind yourself about this conversion rule.

## Supported Variables for Templates

The following table list the supported variables and functions that you can use in templates.

### Supported Variables

Name	Example	Result	Description
1	<code>vm\${1}-example.com</code>	<code>vm172-example.com / 192.168.1.1</code>	The first octet (quad for IPv6) of the discovered asset. Alias for "ip_address_octet1".



Name	Example	Result	Description
2	vm\${2}-example.com	vm41-example.com / 192.168.1.1	The second octet (quad for IPv6) of the discovered asset. Alias for "ip_address_octet2".
3	vm\${3}-example.com	vm13-example.com / 192.168.1.1	The third octet (quad for IPv6) of the discovered asset. Alias for "ip_address_octet3".
4	vm\${4}-example.com	vm9-example.com / 192.168.1.1	The fourth octet (quad for IPv6) of the discovered asset. Alias for "ip_address_octet4".
<i>discovered_name</i>	iface-\${discovered_name}.example.org	iface-example09.example.org / iface-example09	The discovered name of the asset.
<i>ip_address</i>	\${ip_address}.example.org	1.2.3.4.example.org / 1.2.3.4	The IP address of the discovered asset.
<i>ip_address[index]</i>	\${ip_address[7]}.example.com	b.example.com / 2001:db8:acad::1	The IP address octet (quad) substitution. Useful for IPv6 addresses. Throws an error if address have less octets (quads) than specified.
<i>ip_address_octet1</i>	dev-\${ip_address_octet1}.example.com	dev-2.example.com / 2dba::db8::1	The first octet (quad) of the discovered asset.
<i>ip_address_octet2</i>	dev-\${ip_address_octet2}.example.com	dev-d.example.com / 2dba::db8::1	The second octet (quad) of the discovered asset.
<i>ip_address_octet3</i>	dev-\${ip_address_octet3}.example.com	dev-b.example.com / 2dba::db8::1	The third octet (quad) of the discovered asset.
<i>ip_address_octet4</i>	dev-\${ip_address_octet4}.example.com	dev-a.example.com / 2dba::db8::1	The fourth octet (quad) of the discovered asset.

**Supported Functions**

Name	Example	Result	Description
<i>dashed</i>	<code>\${dashed(\${ip_address})}-vm.example.com</code>	1-2-3-4-vm.example.com / 1.2.3.4	Replaces the dot "." and colon ":" symbols with the hyphen symbol "-".
<i>reversed</i>	<code>\${reversed(\${ip_address})}-vm.example.com</code>	4.3.2.1-vm.example.com / 1.2.3.4	Reverts the octets of the IP address. IPv6 address is converted to the dotted representation.
<i>underscored</i>	<code>\${underscored(\${ip_address})}-vm.example.com</code>	1_2_3_4-vm.example.com / 1.2.3.4	Replaces the dot "." and colon ":" symbols with the underscore symbol "_".

### Supported Conversion Parameters

The following table lists the supported magic variables, variables, operators, and functions that you can use to build the formula for the automatic conversion rules.

#### *Magic Variables for Conversion Rules*

Some of the functions or predicates use the following magic variables to calculate the matching results.

Name	Example	Description
<i>ip_address</i>	<code>\${ip_address}</code>	The IP address. Used in <code>is_ipv4</code> , <code>is_ipv6</code> predicates.
<i>mgmt_ip_address</i>	<code>\${mgmt_ip_address}</code>	The management IP address. Used in <code>is_interface</code> predicate.

#### *Supported Variables for Conversion Rules*

This table lists all the variables you can use in the condition syntax.

Name	Discovered by Network Insight	Description
<code>ip_address</code>	Y	Discovered IP address.
<code>mac_address</code>	Y	Discovered MAC address.
<code>duid</code>	Y	DUID associated with the IPv6 address.
<code>netbios_name</code>	Y	Discovered NetBIOS name.
<code>os</code>	Y	OS guessed by network discovery.
<code>method</code>	Y	The method being used for network discovery: FULL, ICMP, NETBIOS, TCP, or CSV.
<code>network_component_type</code>	Y	The type of network component, such as Switch, Router, and others.
<code>network_component_name</code>	Y	The name of the network component.

Name	Discovered by Network Insight	Description
network_component_description	Y	A descriptive string for the network component
network_component_ip	Y	IP Address of the network component.
network_component_port_number	Y	Port number on the network component on which the IP was discovered.
network_component_port_name	Y	Port name on the network component on which the IP was discovered.
network_component_vendor	Y	Vendor name of the network component to which the device is connected.
network_component_model	Y	Model name of the network component to which the device is connected in the vendor terminology.
network_component_port_id	Y	Interface ID of the connected switch/switch-router.
port_vlan_name	Y	Name of the VLAN on the port on the network component.
port_vlan_description	Y	Description of the VLAN on the port on the network component.
port_vlan_number	Y	Number of the VLAN on the port on the network component.
port_speed	Y	Speed settings on the port on the network component: 10M, 100M, 1G, 10G, 100G, or Unknown.
port_duplex	Y	Duplex settings on the port on the network component.
port_status	Y	Status of the port on the network component.
port_link_status	Y	Link Status of the port on the network component.
port_type	Y	Type of interface on the network component to which the device is connected.
open_ports	Y	List of opened ports on the IP address, represented as: "TCP: 21,22,23 UDP: 137,139". Limited to max total of 1000 ports.
last_discovered_timestamp	Y	The timestamp when this data discovered.
first_discovered_timestamp	Y	The timestamp when this IP was first seen by the discovery station.
discovered_name	Y	Name of the IP as seen by the discovery station.
discoverer	Y	Name of the discoverer or Grid member.
device_vendor	Y	Vendor name of the device.

Name	Discovered by Network Insight	Description
vswitch_segment_type	N/A	Type of network segment on which the current virtual machine/vport is connected.
vswitch_tep_ip	N/A	IP address of the virtual tunnel endpoint (VTEP) in the virtual switch.
vswitch_tep_port_group	N/A	Port group of the virtual tunnel endpoint (VTEP) in the virtual switch.
vswitch_tep_vlan	N/A	VLAN of the virtual tunnel endpoint (VTEP) in the virtual switch.
vswitch_tep_dhcp_server	N/A	DHCP server of the virtual tunnel endpoint (VTEP) in the virtual switch.
vswitch_tep_multicast	N/A	Multicast address of the virtual tunnel endpoint (VTEP) in the virtual switch.
vmhost_ip_address	N/A	IP address of the physical node on which the virtual machine is hosted.
vmhost_name	N/A	Name of the physical node on which the virtual machine is hosted.
vmhost_mac_address	N/A	MAC address of the physical node on which the virtual machine is hosted.
vmhost_subnet_cidr	N/A	CIDR subnet of the physical node on which the virtual machine is hosted.
vmhost_nic_names	N/A	List of all physical port names used by the virtual switch on the physical node on which the virtual machine is hosted. Represented as: eth1,eth2,eth3.
vmi_tenant_id	N/A	ID of the tenant to which the virtual machine belongs.
cmp_type	N/A	If the IP is coming from a Cloud environment, the Cloud Management Platform type.
vmi_ip_type	N/A	Discovered IP address type.
vmi_ip_type	N/A	Discovered IP address type.
vmi_private_address	N/A	Private IP address of the virtual machine.
vmi_is_public_address	N/A	Indicates whether the IP address is a public address.
cisco_ise_ssid	N/A	Service Set Identifier.
cisco_ise_security_group	N/A	Name of the security group created in Cisco ISE.
cisco_ise_quarantine_status	N/A	Quarantine status for the IPAddress as coming from Cisco ISE: NONE or QUARANTINE
cisco_ise_endpoint_profile	N/A	Endpoint profile in Cisco ISE.

Name	Discovered by Network Insight	Description
device_type	Y	Type of the device in the vendor terminology.
device_model	Y	Model name of the device in the vendor terminology.
mgmt_ip_address	Y	Management IP address of the device if the device has more than one IP.
device_port_name	Y	System name of the interface with which the IP associates.
device_port_type	Y	Hardware type of the interface with which the IP associates.
is_end_host	Y	Whether this object is an end host or an infrastructure device for the purpose of discovery.
iprg_id	Y	Port Redundant Group ID of this device interface.
iprg_no	Y	Port Redundant Group no of this device interface.
iprg_type	Y	Type of Port Redundant Group
iprg_state	Y	State of this IP address in the group.
vmi_name	N/A	Name of the virtual machine.
vmi_id	N/A	ID of the virtual machine.
vlan_port_group	N/A	Port group to which the virtual machine belongs.
vswitch_name	N/A	Name of the virtual switch.
vswitch_id	N/A	ID of the virtual switch.
vswitch_type	N/A	Type of the virtual switch: standard or distributed: Unknown, Standard, or Distributed
vswitch_ipv6_enabled	N/A	Indicates whether the virtual switch has IPV6 enabled: true or false
vport_name	N/A	Name of the network adapter on the virtual switch connected with the virtual machine.
vport_mac_address	N/A	MAC address of the network adapter on the virtual switch to which the virtual machine is connected.
vport_link_status	N/A	Link status of the network adapter on the virtual switch to which the virtual machine is connected.

Name	Discovered by Network Insight	Description
vport_conf_speed	N/A	Configured speed of the network adapter on the virtual switch to which the virtual machine is connected. Unit is Kib.
vport_conf_mode	N/A	Configured mode of the network adapter on the virtual switch to which the virtual machine is connected: Unknown, Full-duplex, or Half-duplex
vport_speed	N/A	Actual speed of the network adapter on the virtual switch to which the virtual machine is connected. Unit is Kib.

### Supported Operators for Conversion Rules

Operators always result in boolean value: true or false. Therefore, you can use them only in logical expressions.

Name	Left Value (lvalue)	Right Value (rvalue)	Example	Description
LIKE	variable	string (regular expression in extended format)	<code>\${discovered_name} like "[vV]m-[0-9]+.devnet.org"</code>	Evaluates as true if the lvalue variable matches the given regular expression rvalue; otherwise false
==	variable	string	<code>\${ip_address} == "167.45.13.29"</code>	Evaluates to true if the lvalue variable equals to rvalue string literal, false otherwise
!=	variable	string	<code>\${mac_address} != "00:50:56:00:00:01"</code>	Evaluates to true if the lvalue variable is not equal to rvalue string literal, false otherwise

### Supported Functions or Predicates for Conversion Rules

Predicates accept either none or one argument. Depending on the predicate, it could accept both variables and strings or only one of them. The predicate can be compared only to boolean value: true or false.

Name	Argument Type	Example	Description
is_interface	N/A	<code>\${is_interface} == true</code>	Check discovered data in an interface. It validates the mgmt_ip_address variable.
is_ipv4	N/A	<code>\${is_ipv4} == true</code>	Check to see if the variable ip_address is an IPv4 address.
is_ipv6	N/A	<code>\${is_ipv6} == false</code>	Check to see if the variable ip_address is an IPv6 address.
is_belongs_to	string	<code>\$(ip_belongs_to("10.0.0.0/8")) == false</code>	Check to see if the ip_address variable belongs to the given IPv4 or IPv6 network range.

Name	Argument Type	Example	Description
is_fqdn	variable	<pre> \${is_fqdn( {discovered_name})} == true </pre>	Check to see if the given variable is an FQDN.

## Accessing Detailed Device Information

Clicking on any device's name in the **Name** column, you open a set of tabs revealing information about the selected device:

- The **Interfaces** tab (see [Viewing Interface Information for Discovered Devices](#) below)
- The **Networks** tab (see [Viewing Networks Associated with Discovered Devices](#) below)
- The **IP Addresses** tab (see [Viewing IP Addresses Associated with Discovered Devices](#) )
- The **Assets** tab (see [Viewing Assets Associated with Discovered Devices](#) below)
- The **Components** tab (see [Viewing Components of Discovered and Managed Devices](#) below)

### Viewing Interface Information for Discovered Devices

This panel lists all discovered interfaces associated with the selected device. Interfaces are detected whether they are loopbacks, unnumbered, or numbered with one or more IP addresses. Interfaces may be listed for either managed or unmanaged devices.

1. From the **Data Management** tab, select the **Devices** tab. The Devices Home page displays a list of all devices currently found and catalogued by discovery.
2. Click the **Action** icon




for a chosen device and choose **Interfaces** from the drop-down menu, or simply click the device name to display the Interfaces list. Click **Devices Home** to return to the main **Devices** page.

This panel displays the following information for each interface. Note that some data may appear for some device types and not for others.

- **Name:** The name of the interface (usually a switched interface) associated with the discovered device.
- **Reservation:** Indicates whether the port has been reserved by NIOS as part of a Port Control operation.
- **IP Address:** Detected IPv4 or IPv6 address of the interface.
- **VRF Name:** The name of the VRF associated with the interface, if applicable.
- **Network View:** The name of the network view to which the VRF instance belongs, if applicable. If there is only one network view in the Grid, which is the default view, the **Network View** column is hidden by default.
- **VRF Description:** The description about the VRF instance, if applicable.
- **VRF RD:** The route distinguisher associated with the VRF instance, if applicable.
- **MAC Address:** The hardware address associated with the interface.
- **Description:** Port description associated with the interface, such as **ge-0/0/5** or **FastEthernet0/13**.
- **VLAN ID/VLAN Name:** The data VLAN identifier and VLAN name that is bound to the interface, if applicable.
- **Port Speed:** Interface speed, in Mbps.
- **Port Type:** Type of interface as detected by NIOS Discovery. Examples include **ethernet-csmacd**, **propPointToPoint Serial**, **I2vlan**, **tunnel**, and others.
- **Admin Status:** Shows whether the interface is administratively Up or administratively Down.
- **Operating Status:** Shows whether the interface is operationally Up or operationally Down.
- **Trunk Status:** Where applicable, shows the trunking status of the interface.
- **Link Aggregation:** Shows if the interface is part of a Link Aggregation Group, also known as port channel.
- **Aggregated Interface:** Shows the port channel name, if the interface is part of a port channel or virtual port channel.
- **vPC Peer:** This field is applicable to Cisco devices with virtual port channel. It shows the name of the second aggregated interface of the port channel. The aggregated interface and vPC peer values are identical as aggregated interfaces on a vPC device must have identical names. Clicking the link in the **vPC Peer** column

takes you to the **Interfaces** tab of the device to which the peer interface belongs. Also see [IP List Neighbor Information](#) for vPC data displayed in the **Attached Device Port Name** field for an end device.

- **Status:** Shows whether the interface is Used or Unused.
- **IPAM Type:** The object type that is associated with the IP address for the interface. Possible values can be Lease, IPv4 DHCP Range or Fixed Address.
- **Usage:** Indicates whether NIOS has configured the IP address for DNS or DHCP.
- **Managed:** Shows whether the interface is managed under IPAM, by being associated with a managed IPAM object such as a Fixed Address. Check the IPAM Type field for related information.
- **Reservation:** Indicates whether the interface has a port reservation bound to it. For information, see [Creating Port Reservations for IPAM Objects](#).
- **Capabilities:** Describes the capabilities of each interface in the selected device. Hover the mouse over each entry to view the complete listing. For information, see [Determining Interface Capabilities](#) below.
- **Site:** This is a predefined extensible attribute.

You can also click the Action icon  next to an interface name and select one of the following to perform the specified task:

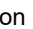
- **Edit:** This option is enabled if the network is in managed state. This opens the network editor.
- **Show Assets:** This option is only available for switched Ethernet interfaces with no IP Address. This opens the Assets page for the selected device, and shows a list of end host devices or neighboring linked to the interface. Network Insight filters the asset list for the device by the interface name.
- **Show Multiple IP Addresses:** Opens the IP Addresses page specifically for the interface, listing all IPv4 and IPv6 addresses associated with the interface. This option appears only if the interface has IP addresses.
- **Convert:** Convert a network in the unmanaged state to be managed under IPAM and (optionally) DHCP. Unlike devices and interfaces, you do not assign objects such as fixed addresses or PTR records to a managed network. Conversion enables a network to be fully manageable under IPAM and DHCP. For more information, see [Converting Unmanaged Networks to Managed Status](#).
- **IPAM Networks:** Choosing this option lists all IPAM networks associated with the current interface.
- **Device Details:** A basic list of information about the chosen device, including the IP address by which the device is discovered, operational status, IPAM Type (whether the device is Managed or Unmanaged), the Device Type and the number of Interfaces.

You can also do the following:

- Modify some of the data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes or click **Cancel** to exit. Note that some fields are read only. You can modify the following fields in this table: **VLAN ID/VLAN Name**, **Admin Status** and **Description**.
- Sort the data in ascending or descending order by column.
- Select an interface checkbox and click the Edit icon to manage device properties.
- Click the Export icon to export the list of discovered devices to a .csv file.
- Click the Print icon to print the list of discovered devices.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of a device name in the **Go to** field and select the device from the possible matches.
- Create a quick filter to save frequently used filter criteria.

## Determining Interface Capabilities

One key piece of information that Network Insight compiles from discovery involves an assessment of interface *capabilities*. The Interfaces table provides a hidden column titled **Capabilities** with a description of the relevant capabilities of each interface in the discovered device.

1. From the **Data Management** tab, select the **Devices** tab. The Devices Home page displays a list of all devices currently found and catalogued by discovery.
2. Click the Action icon  for a chosen device and choose **Interfaces** from the popup menu.
3. Click the right end of a column header and choose **Columns** → **Edit Columns** from the drop-down menu.
4. Select the **Capabilities** checkbox and click **Apply**.

The text listing for the Capabilities field may be too long to display in the Interfaces table. Hover the mouse over any table row to display the complete entry for the **Capabilities** field:



Capabilities information for the selected interface

USAGE	MANAGED	RESERVATION	CAPABILITIES
			Provision IPv4 Network = No: Device type (Wireless Controller) not supported for network provisioning...
			Provision IPv4 Network = No: Device type (Wireless Controller) not supported for network provisioning...
			Provision IPv4 Network = No: Device type (Wireless Controller) not supported for network provisioning...
			Provision IPv4 Network = No: Device type (Wireless Controller) not supported for network provisioning...
	No		Provision IPv4 Network = No: Device type (Wireless Controller) not supported for network provisioning...

Critical values in the Capabilities field include the following:

- Provision IPv4 Network – Yes or No
- Provision IPv6 Network – Yes or No
- De-Provision IPv4 Network – Yes or No
- De-Provision IPv6 Network – Yes or No
- Interface has no IP Address – Appears only if no IP address is defined for the interface
- Edit Data VLAN – Yes or No
- Edit Voice VLAN – Yes or No
- Edit Admin Status – Yes or No
- Edit Description – Yes or No

The values you see in this field provide notification when you are unable to set certain values for any interface. When discovery queries devices for their device type, OS version support and other factors, Network Insight compiles the information into table formats that help you to determine issues or exceptions in the network.

For example, loopback interfaces are "always up" and hence do not support an Admin Status configuration and show **Edit Admin Status = No**. If an interface does not have an IP address, it cannot be de-provisioned and hence shows a **De-provision Network = No**. That same device may not support IP addresses or VLANs because it is only a L2 Ethernet switch, and shows **Edit VLAN = No** and **Provision IPv4 Network = No**, among others.

If a device shows a **Provision IPv4 Network = Yes** and **Provision IPv6 Network = No**, it indicates that the device supports only IPv4, perhaps due to OS software version.

Other cases may involve the following reasons:

- **Edit Admin Status = No:**
  - Device not supported
  - Vendor not supported
  - Model not supported
  - Port does not support Admin settings (Loopback, Virtual)
- **Edit VLAN = No:**
  - Device not supported
  - Device type not supported for VLAN assignment (Router, L2 switch)
  - Vendor not supported
  - model not supported
- **Edit Voice VLAN = No:**
  - Device not supported
  - Vendor not supported (Cisco is the only supported device type)
  - Device not licensed
  - Model not supported
- **Provision IPv4 Network=No/Provision IPv6 Network=No:**
  - Device not supported
  - Vendor not supported for network provisioning
- **De-Provision IPv4 Network=No/De-Provision IPv6 Network=No:**
  - Device not supported
  - Vendor not supported for network provisioning
  - Device type not supported for network provisioning (Router)
- **Edit Description=No**
  - Device not supported

- Vendor not supported
- Model not supported

## Viewing Networks Associated with Discovered Devices

To view all discovered networks associated with the selected device, complete the following:


1. From the **Data Management** tab, select the **Devices** tab. The Devices Home page displays a list of all devices currently found and catalogued by discovery.
2. Click the device name.
3. Click the **Networks** tab.
 

Grid Manager displays all networks to which the chosen device connects. You see the following types of networks based on their managed or unmanaged status:

  - **Unmanaged:** These networks are displayed in yellow rows with the value of **No** in the Managed column. Shows that the network is not managed under IPAM, but enough network information is catalogued so that the network can be converted to managed status. You can provision these networks onto devices.
  - **Managed:** These networks are displayed in white rows with the value of **Yes** in the Managed column. These networks are currently managed under IPAM, converted to an IPAM network. You can provision and de-provision managed networks.
  - **The so called "non-NIOS networks":** These networks are displayed in grey rows with a blank value in the Managed column. Indicates that the network is discovered but available network information is not sufficient to identify and catalog the network in IPAM at the present time. If you encounter such networks and you want them to appear in IPAM, do the following:
    - Check the admin or operation status of the corresponding interface. It should not be disconnected physically or disabled by administrator.
    - Ensure that the prefix length for the network is other than /32 (ipv4) or /128 (ipv6). Network Insight treats named prefixes as a route to a specific device rather than a subnet, therefore it does not create such network in IPAM.
    - Ensure that the route for this interface is configured correctly.
    - Check that the route is a direct or local one based on the routing table and is not learned from a remote source via BGP, OSPF and so on (i.e., indirect next hop), nor comes from a static route using the netmgmt protocol.
    - If the network is within a VRF, ensure that the VRF is mapped to the correct network view. VRF mapping is required in this case for the network to appear in IPAM. After the VRF is correctly mapped, the network turns from non-NIOS to unmanaged, or managed if the network is already present in IPAM.

Grid Manager displays the following information for each network found on the selected device, if applicable:

- **Network:** The network IPv4 or IPv6 address.
- **VRF Name:** The name of the VRF associated with the interface, if applicable.
- **Network View:** The name of the network view to which the VRF instance belongs, if applicable. If there is only one network view in the Grid, which is the default view, the **Network View** column is hidden by default.
- **VRF Description:** The description about the VRF instance, if applicable.
- **VRF RD:** The route distinguisher associated with the VRF instance, if applicable.
- **Comment:** Any information entered by admins about the network.
- **Managed:** Shows values of Yes or No for managed status.

Using the Action icon , you can perform the following tasks in the Networks page:

- **Show IPAM Network:** Opens the IPAM IP MAP that illustrates the IP states for all IPs in the network.
- **De-provision Network:** Available for managed networks that are provisioned and active on a device. Allows you to de-provision (delete) the selected IPAM network from all devices connected to the selected network. See [Provisioning and De-Provisioning Networks](#).
- **Edit:** Opens the network editor for the selected network. This option is enabled if the network is in managed status.
- **Delete:** Select **Delete** to delete the network now or select **Schedule Deletion** to schedule the deletion at a later time. Note that the deletion function allows you to de-provision the actual network from the device. By default, when you choose **Delete** or **Scheduled Delete**, the network is de-provisioned from all interfaces listed in the panel. Exercise caution when using this feature!
- **Extensible Attributes:** Provides access to the extensible attribute settings for the selected network.

- **Permissions:** Provides access to admin permissions settings for the selected network. This option is enabled if the network is in managed status.
- **Convert:** Converts unmanaged network to a managed network in NIOS. All discovered networks on each device are automatically listed as **Unmanaged** after a discovery. This means that the discovered network, though visible, does not have its identities resolved by NIOS, nor are its IP address managed through IPAM or leased through DHCP. After converting the unmanaged network to managed status, Grid Manager uses the discovered router IP address to populate the same value under subsequent DHCP configurations for the network. You can also select an unmanaged network and convert it to managed status by clicking **Convert** from the Toolbar.
- **Device Details:** Provides information about the device to which the selected network belongs. The list includes information such as the **IP Address** and **Device Type** for the device, and in the **IPAM Type** field whether the device itself is a managed or unmanaged object in NIOS. It also provides the following status counters for the device:
  - **Administrative Up - Operational Up:** The number of ports that are fully up and passing traffic
  - **Administrative Up - Operations Down:** The number of ports that are administratively up, but have some kind of connectivity issues.
  - **Administrative Down - Operational Down:** The number of ports that are administratively taken down.

The horizontal navigation bar and the Toolbar also provide the following functions:

- **Provision Network**

: Available for managed networks and for unmanaged networks that are recognized by IPAM. For information, see [Provisioning and De-Provisioning Networks](#). Clicking this icon opens the Provision Network feature, allowing you to provision the network onto the actual device by selecting a device interface, and enabling DHCP Forwarding and/or assigning a VLAN. Grid Manager creates a new port control task under Task Manager, and you can choose the interface on which the network is provisioned, along with VLAN configuration and other settings.

- **De-Provision Network**

: Available for discovered networks that are not visible under IPAM. A dialog box appears summarizing the task.

- **Show Active Users:** For Microsoft Management only. Displays the *Active Users* dialog box. You can view all the active users on the Active Directory domain for the selected device. For more information, see [Viewing Active Network Users](#).

## Modifying Networks

Grid Manager enables the user to edit select DHCP configurations, including the following:

**IPv4 DHCP Options/IPv6 DHCP Options:** DHCP options provide specific configuration and service information to DHCP clients. For more information, see [About IPv4 DHCP Options](#).

**DHCP Forwarding:** Enables routers connecting multiple networks to act as a silent DHCP relay and forward DHCP packets between them. The **DHCP Forwarding** page lists the interfaces on the currently selected network on which DHCP Forwarding is enabled. If more than one device on the selected network also enables DHCP Forwarding, they also appear here. DHCP Forwarding configuration involves simply enabling or disabling the service for a network endpoint on the device. In order for DHCP forwarding to work, you must restart the DHCP service on the Grid member that is serving the network; If you run DHCP service on both LAN1 and LAN2 of the Grid member, then both addresses are written to the device.


1. From the **Data Management** tab, select the **Devices** tab.
2. Click the **Name** link for the device you want to inspect.
3. Click the **Networks** tab.
4. Click the Action icon

for a network in the table, and choose **Edit**. This feature is enabled only for networks that are managed under IPAM.

5. Click the **DHCP Forwarding** tab.
6. Select the checkbox for any listed instance and do the following if necessary:
  - Click **Configure**. Grid Manager queries you to confirm that DHCP Forwarding are configured on the selected network (**A task will be created to configure DHCP forwarding for this network on these devices: <device\_name>. You can view the execution log for the task in the Task Manager to see the results**).
  - Click **Delete** to remove the selected DHCP Forwarding instance from the network.

## DHCP Forwarding confirmations

DHCP Forwarding Configuration ✕

 A task will be created to configure DHCP forwarding for this network on these devices: WS-C3750X-24.inca.infoblox.com. You can view the execution log for the task in the Task Manager to see the results.

Do you want to Continue?

1. Click **Yes** to confirm the activating or deletion of DHCP forwarding on the selection, or **No** to reject the change.
2. Click **Save & Close**.

## Viewing IP Addresses Associated with Discovered Devices

You can view the complete list of discovered IP addresses bound to all interfaces for any device, discovered and managed devices alike.



### Note

One useful trick for interfaces is to pick out an interface from the **Interfaces** page that has multiple IPs and open the **IPAddresses** tab; or sort the IP addresses table by its **IPAddress** column, and locate the interface name that bears multiple IPs. Frequently, an interface with multiple addresses can have IPv4 and IPv6 addresses bound to it. Loopbacks are another example.

1. From the **Data Management** tab, select the **Devices** tab. The Devices Home page displays a list of all devices currently found and catalogued by discovery.
2. Click the Action icon

for a chosen device and choose **Interfaces** from the drop-down menu, or simply click the device name to display the Interfaces list. Click **Devices Home** to return to the main **Devices** page.

3. Click the **IP Addresses** tab. Grid Manager displays all IP addresses associated with the chosen device. Grid Manager displays the following information for each IP address:

- **IP Address:** The IP address for each discovered interface as managed by NIOS and IPAM. The table supports IPv4 and IPv6 values. Each IP address is a link to the home IPAM page for the interface. If an IP address does appear but is not a link, this indicates the discovered IP is not recognized under IPAM.
- **VRF Name:** The name of the VRF associated with the interface, if applicable.
- **Network View:** The name of the network view to which the VRF instance belongs, if applicable. If there is only one network view in the Grid, which is the default view, the **Network View** column is hidden by default.
- **VRF Description:** The description about the VRF instance, if applicable.
- **VRF RD:** The route distinguisher associated with the VRF instance, if applicable.
- **Interface Name:** The name of the interface (usually a switched interface) associated with the discovered device.
- **MAC Address:** The hardware MAC address associated with the interface.

23	198612146
1304	16446202
1313	16446213
400	400
400	132
800a8...	

- **VLAN Name/VLAN ID:** The data VLAN name and VLAN identifier to which the interface is bound, if applicable. In most cases, you see both the VLAN name and the VLAN ID as two values in the same field. Multiple VLAN entries may be present for an interface or IP Address. Some interfaces may have a large number of associated VLANs. By default, Network Insight does not automatically show all of them, instead providing a **Show all...** link for reference within the table cell. All VLAN ID/VLAN name values appear within the table cell, with a **Hide...** link provided to shorten the list back to original length.
- **AdminStatus:** Lists whether the interface is administratively Up or administratively Down.
- **OperationStatus:** The operational status of the interface (operationally Up or operationally Down).

- **Managed:** Indicates whether or not the IP Address is managed by Grid Manager. If the IP address is unmanaged, you will be able to Convert the IP address to an Object that is managed by Grid Manager.
  - **Site:** This is a predefined extensible attribute. Extensible attributes may also appear in this table.
4. Click the IP Address link for any interface to open the Related Objects page for the chosen port. Click the Action icon


next to an IP address and select one of the following to perform the specified task. Note that some of these actions are not applicable to the IP address.

- **Edit Interface:** Opens the interface general settings page. You can view and modify basic interface settings such as **Admin Status** (on the **General** page), **Data VLAN** and **Voice VLAN** (on the **VLAN** page), and add or modify extensible attributes.
- **Convert:** Depending on the address type and its IPAM status, you may be able to convert the selected IP to a Host Record, A Record, PTR Record or a Fixed Address. Otherwise, Grid Manager shows **This object cannot be converted**. You can also perform the same action by selecting an IP address checkbox and clicking **Convert** from the Toolbar.
- **Device Details:** Provides information about the device to which the selected IP address belongs. The list includes information such as the **IP Address** and **Device Type** for the device, and in the **IPAM Type** field whether the device itself is a managed or unmanaged object in NIOS. It also provides the following status counters for the device:
  - **Total Available Interfaces:** The total number of interfaces associated with the device.
  - **Administrative Up - Operational Up:** The number of ports that are fully up and passing traffic
  - **Administrative Up - Operations Down:** The number of ports that are administratively up, but have some kind of connectivity issues.
  - **Administrative Down - Operational Down:** The number of ports that are administratively taken down.

## Viewing Router Redundancy Information

Some discovered devices may support router redundancy. After discovery, some IP addresses are indicated as a VIP (virtual IP) in the device's IP Addresses page with router redundancy. In the **IP List** page, **Router Redundancy** in the Action icon menu lists the IP addresses associated with the VIP. For each IP, there are various menu items as shown in the below figure.

*Virtual IPs and discovered redundancy information*



- **Active:** lists the active interface in the redundancy pair;
- **VIP:** The Virtual IP for the router redundancy pair;
- **Standby:** The standby IP interface for the router redundancy.

Discovery of all three IP components of the Router Redundancy instance also provides related information for all three IP entities:

- **Show IPAM Address:** opens the IPAM page to the listed IP address;
- **VIP:** Opens the virtual interface in the host device's Interfaces page;
- **Device name:** The third item lists the device name of the router for each of the three IP address entities comprising the redundancy instance. The currently active router is identified with the **Active** and **VIP** objects; the second **Standby** router is identified with the **Standby** IP address.


## Viewing Assets Associated with Discovered Devices

During discovery, Network Insight classifies end hosts and any other devices connected to switchport interfaces as "Assets" directly associated with each discovered interface. On the device level, the Assets page shows all network devices reachable by the selected network device, including switchports supporting end hosts. In practice, most Asset tables show end hosts and devices that populate Ethernet network segments.

The Assets table lists all managed end hosts and application servers detected through discovery and identity resolution by Grid Manager, that are connected to each network infrastructure device. The records listed in this table date from the


Last Seen discovery time stamp of each end host or other device. In many cases, you see neighbor devices to the current device appearing on this page.

To view assets associated with discovered devices:

1. From the **Data Management** tab, select the **Devices** tab. The Devices Home page displays a list of all devices currently found and catalogued by discovery.
2. Click the Action icon  


for a chosen device and choose **Interfaces** from the drop-down menu, or simply click the device name to display the Interfaces list. Click **Devices Home** to return to the main **Devices** page.

3. Click the **Assets** tab. Grid Manager displays all assets associated with the chosen device. Note the list of assets at this level may include devices that are trunked to the current device, including end-user host computers or routers and switch-routers neighboring the currently selected device. Grid Manager displays the following information for each asset:
  - **Name:** The asset name on the network as discovered by Grid Manager. If the name is that for another infrastructure device, you may click on it to see its associated assets.
  - **Interface Name:** The name of the interface (usually a switched interface) associated with the discovered device.
  - **VRF Name:** The name of the VRF associated with the interface, if applicable.
  - **Network View:** The name of the network view to which the VRF instance belongs, if applicable. If there is only one network view in the Grid, which is the default view, the **Network View** column is hidden by default.
  - **VRF Description:** The description about the VRF instance, if applicable.
  - **VRF RD:** The route distinguisher associated with the VRF instance, if applicable.
  - **IP Address:** The IP Address for each discovered asset as managed by NIOS and IPAM. The IP address is a link to the home IPAM page for the interface.
  - **Type:** The type of device. Infrastructure devices such as routers and switches may also be categorized as an Asset.
  - **Username:** The User Name for the asset, as defined from the host's DHCP lease.
  - **Asset MAC Address:** The hardware MAC address associated with the asset.
  - **VLAN Name/VLAN ID:** The VLAN identifier from which the asset is reachable.
  - **Admin Status:** The administrative status (Up or Down) of the management port that identifies the asset device.
  - **Operation Status:** Normally reads **Up** or **Down**. Asset records may appear as **Down** because they are disconnected from the network or being rebooted.
  - **Managed:** Indicates whether the asset is managed by NIOS.
  - **Reservation:** Indicates whether the interface has a port reservation bound to it.
  - **Capabilities:** Describes the capabilities of each interface in the selected device. Hover the mouse over each entry to view the complete listing.
  - **Site:** This is a predefined extensible attribute.

Using the Action icon , you can perform the following tasks in the IP Addresses page:

- **Edit IPAM Object:** For managed objects, this opens the editor for the object so you can modify its properties.
- **Edit Interface:** Opens the interface general settings page. You can view and modify basic interface settings such as **Admin Status** (on the General page), **Data VLAN** and **Voice VLAN** (on the VLAN page), and add or modify extensible attributes.
- **Show IPAM IP Address:** Shows discovered data, related objects, and audit history of the selected asset. This option is disabled for devices that have a management IP that is not part of an IPAM network. Discovered data may or may not appear for the asset, depending on the device type.
- **Convert:** Depending on the asset type and its IPAM status, you may be able to convert the selected asset to a Host Record, A Record, PTR Record or a Fixed Address. Otherwise, Grid Manager shows **This object cannot be converted**. You can also perform the same action by selecting an IP address checkbox and clicking **Convert** from the Toolbar.
- **Device Details:** Provides information about the device to which the selected IP address belongs. The list includes information such as the **IP Address** and **Device Type** for the device, and in the **IPAM Type** field whether the device itself is a managed or unmanaged object in NIOS. It also provides the following status counters for the device:
  - **Total Available Interfaces:** The total number of interfaces associated with the device.
  - **Administrative Up - Operational Up:** The number of ports that are fully up and passing traffic
  - **Administrative Up - Operations Down:** The number of ports that are administratively up, but have some kind of connectivity issues.

- **Administrative Down - Operational Down:** The number of ports that are administratively taken down.

### Viewing Assets Associated with an Interface

You can narrow down an Asset list to individual interfaces on any managed network device. On the interface level, the **Assets** page shows all devices associated with the chosen interface, including switchports supporting many end hosts. In practice, most Asset tables show end hosts and devices that populate Ethernet network segments.

1. From the **Data Management** tab, select the **Devices** tab.
2. Click the **Name** link for the device you want to inspect.
3. Click the Action icon

for an interface in the table, and choose **Show Assets**. (Applies only to switched interfaces that do not have an IP address.)

Values listed in the Assets table include the following:

- **Name:** The asset's name on the network as discovered by Grid Manager. If the Name is that for another infrastructure device, you may click on it to see its associated Assets.
- **Interface Name:** The name of the interface (typically a switched interface) associated with the asset (by which the asset was discovered).
- **IP Address:** The IP Address for each discovered end host as managed by IPAM. The IP address is a link to the home IPAM page for the interface.
- **MAC Address:** The hardware MAC address associated with the asset.
- **VLAN ID/VLAN Name:** The VLAN identifier from which the asset is reachable.
- **Operation Status:** Normally reads **Up** or **Down**. Asset records may appear as **Down** because they are disconnected from the network or being rebooted.

In the Interfaces page, if you select an interface for a switch that is only connected to a neighboring switch, router, or switch-router, and then choose **Show Assets**, the Assets page displays only the neighboring device that is reachable from the chosen port.

### Viewing Components of Discovered and Managed Devices

Network Insight provides a table of hardware components for each discovered/managed network device. Elements listed in the Components table include the following:

- **Name:** The discovered name of the device component.
- **Description:** The description string associated with the component. In many cases, this value is the same as the Name field.
- **Class:** Type of component. Possible values include, but are not limited to, Port, Power Supply, Fan, Module, Stack, RoutingEngine, powerEntryModule, Chassis, and more, based upon the collected SNMP data.
- **Serial Number:** The discovered vendor serial number for the component.
- **Model:** Describes the model number or model name of the component, based upon collected SNMP data.



#### Note

- Note To view information about the device, including its IPAM Type and the operating status of its ports, click the **Action** icon



for a chosen device and choose **Device Details** in the drop-down menu.



## Adding Discovery Device Support



### Note

Adding Device Support Bundles, viewing and deleting them requires Superuser permissions.


Infoblox frequently provides support files for additional network devices that may not previously be supported by discovery, and updates to support new operating system versions of existing devices. To add device support updates:

1. From the **Grid** tab, select the **Device Support** tab.
2. Expand the Toolbar and click **Add**.
3. Click **Select** and navigate to the file you want to upload.
4. Select the file, and then click **Upload**.

The Device Support table shows its installed library of files with the following data points:

- **Name:** The descriptive device name for the device support file.
- **Version:** The version of the currently active device support file.
- **Author:** The developer of the device support file.
- **Type:** The Type column lists two types of Device Support files: the System type indicates a support bundle that is installed with the NIOS/Grid Manager system. The Downloaded type indicates device support bundles that are installed by the administrator.

System bundles are read-only and cannot be removed or overwritten by administrators.

You may remove custom support bundles that you have installed on the Device Support Page. To do so, click the Action icon  for a chosen device and choose **Delete** from the popup menu.

## Excluding IP Addresses from Discovery



### Note

You use the IPv4 Network or IPv6 Network editors to exclude IP addresses or ranges of IP addresses from discovery within the specified network. For more information, see [Disabling Discovery for a Network](#).

Host records, fixed address and IPv4 reservations can be excluded from discovery. You may also exclude an IP address or a range of IPs within a network from discovery.



### Note

You may create a network and choose not to discover it at that time, by disabling both **Enable Immediate Discovery** and **Enable Discovery**. If you disable the **Enable Discovery** checkbox, the network will never be discovered unless you change the setting again at a later time. Conversely, you can explicitly exclude specific IPs or IP ranges from discovery. Discovery will never take place on these IPs unless the admin specifically changes their exclusion setting.

Administrators can specify IPv4 and/or IPv6 addresses that must be immediately discovered by the appliance. Some devices may need exclusion because they do not support SNMP, or for other organizational reasons. Devices matching IP addresses selected for immediate discovery (using the **Discover Now** feature described in [Using Discover Now to Discover an Existing Object](#)) are given one-time priority over other discovered devices, for data collection and counting toward any device found matching the license limits. A device specified through an IP address can also be excluded from discovery or management.




## Quick Exclusion of IPs from Discovery

You can use the IPAM IP Map or IP List page to quickly exclude IP addresses and selected ranges of IP addresses from discovery. For example, you may have Infoblox appliances or routers that provide the gateway to networks that are already managed by Grid Manager, or devices on which you do not wish to have discovery operations take place.

1. From the **Data Management** tab, select the **IPAM** tab. The IPAM Home page appears.
2. Click on any network or network container in the list. The **IP Map** appears for the selected network.  
Note that you may also use the **List** page for the selected network to exclude IPs or selected ranges of IPs. However, you have to page through or search through the pages comprising the list view to locate the IPs you want to exclude. (If you know the IP address value in the **List** view but it does not appear on the page, enter it in the **Go to** field to search for the IP.) The **IP Map** view allows you to view every IP address in a selected network, such as a /24 prefix.
3. Select one or more IPs in the map. SHIFT+click to select a series of contiguous IPs. CTRL+click to select non-contiguous IPs.
4. Expand the Toolbar and click **Exclusion → Enable Exclusion**. The selected IP addresses are excluded from any discovery actions.



### Note

You can click the Action icon  for any List record and choose **Exclusion→EnableExclusion**.

To locate an IP to exclude within a network container:

1. From the **Data Management** tab, select the **IPAM** tab. The IPAM Home page appears.
2. Select the network container by clicking it. The IPAM Home page changes to display the List page, showing the list of networks within the container.
3. Click the network that has the IP in its space that you wish to exclude.
4. Select the network IP address from the **List** table. If you know the IP address value but it does not appear on the page, enter it in the **Go to** field to search for the IP.
5. Expand the Toolbar and choose **Exclusion → Enable Exclusion**.

A parent network container may exclude IPs, and you may add and remove discovery exclusions within network containers. You can drill down to the child networks in the **Net Map** view to perform exclusions on IPs. For example, consider a /16 network container that has a number of smaller /24 child networks within it. Right-clicking on any child network in the network container and choosing **Edit** from the popup menu, opens the editor with its **Discovery Exclusions** tab, where you can perform exclusions within the child network.

## Creating a New IPAM Object and Excluding it from Discovery

You can create a new object from the main **Data Management → IPAM** or **Data Management → DHCP** pages.

1. From the **Data Management** tab, select the **IPAM** or **DHCP** tab.
  - a. For IPAM, the **IPAM** home page appears, listing all networks reachable by IPAM.
  - b. For DHCP, the home **Networks** page appears, showing all networks reachable by DHCP.



### Note

Note the network lists between IPAM and DHCP will likely differ, because networks can be set to be Disabled from DHCP. IPAM provides a complete list of all networks configured or discovered by Grid Manager.

2. To select a network, do one of the following:
  - On the **IPAM** page, select the network.
  - Click **DHCP→Networks** and then select the network.

3. On the Toolbar, expand the **Add** drop-down list and then select the object type from the drop-down menu, such as **Host, Fixed Address** → **IPv4**, or another object type.

4. In the second Wizard step, click **Next Available IP** to obtain the next available IP address in the chosen network. For more information about obtaining the next available IP address, see [About the Next Available Network or IP](#).



Note

For adding a host record, the first step in the Add Host Wizard requires adding an IP address.

5. If the network of the IP address is served by a Grid member, Grid Manager displays the **Assign IP Address by** section, with its **MAC Address**, **DHCP Client Identifier** and **DHCP Relay Agent** settings. Select the different options as needed to define a fixed IP Address or another object.

6. Click **Next** to continue to the DHCP Options page in the wizard.



Note

For more information about DHCP Options configuration, see [About the Next Available Network or IP](#).

7. If you do not wish to configure DHCP Options for this Fixed IP, click **Next** to go to the following Wizard step.

8. (*Optional*) In the **Device Information** page, select the Device Type and the Device Vendor, or, if you do not wish to configure device settings for the current object, click **Next** to go to the next Wizard step, for defining discovery settings.



Note

For more information about Device Information settings, see [Creating Port Reservations for IPAM Objects](#).

9. Choose from the options on the Step 5 Wizard page:

- Check **Exclude from Network Discovery** to prevent the object from being probed by discovery.
- (Enabled by default) If you want immediate discovery of the current Fixed IP, select the **Enable Immediate Discovery** checkbox.  
You may disable both checkboxes. Doing so does not disable discovery for the current object—discovery is simply performed by Network Insight on its own internal timetable.
- To override SNMP credentials for either SNMPv1/v2 or for SNMPv3 for the current Fixed IP, select the **Override Credentials** checkbox.  
You can enter both a SNMPv1/v2 Read community string and an SNMPv3 credential, or enter only the single type you need. Each can be selected and edited in turn.
  - Select **SNMPv1/v2** and enter the **Read** community string;  
—or—
  - Select **SNMPv3** and enter the device admin account name, and the **Auth Protocol**, **Auth Password**, **Privacy Protocol** and **Privacy Password** values where necessary.
- You can apply a set of CLI Credentials to the specific fixed address object. To override the inherited CLI credentials from the Grid, select the **Override CLI Credentials** checkbox. You can enter the admin user **Name** and **Password** values and, if necessary, an **Enable Password**. The values you enter here are specific only to the current object.



#### Note

Clicking **Schedule for Later** is a navigational button to allow you to skip quickly to the scheduling step in the Wizard. You can return at any time to complete remaining Wizard steps to finish creating the object. You may click the **Schedule for Later** button at any time in the Wizard process. For more information, see [Scheduling New IPAM/DHCP Objects and Associated Port Configurations](#).

10. (Optional) Click **Next** to go to the sixth Wizard step, which governs **Reserve Port** and **Configure Port** settings. (For more information, see [Creating Port Reservations for IPAM Objects](#).) Click **Next** when finished with settings.

11. If necessary, add or apply any extensible attributes necessary for the new record. Click **Next**.

12. The final Wizard page governs scheduling of the object creation task and the optional port reservation task. Click **Save & Close**.

## Excluding IP Addresses in Grid Manager

You may exclude IP addresses from discovery from within a number of different contexts in Grid Manager. Under IPAM, you can exclude in the IP Map and IP List pages. The IP List page provides an Exclude data column that directly shows the exclusion status for all IPs in the selected network. Various objects, such as host records, IPv4 and IPv6 fixed addresses, and IPv4 reservations, may be excluded from discovery.

You may view excluded IP addresses in the IP List page or in the network editor's **Exclusions** tab.

## Disabling Discovery for a Network

You go to the DHCP feature under **Data Management** to disable discovery for a network. Disabling discovery for a network differs from discovery blackouts; disabling discovery for a network simply ensures that discovery never takes place on the chosen network.

To disable discovery for an IP network:

1. Select a managed network from one of the following locations:
  - a. **Data Management** → **IPAM** → list view
  - b. **Data Management** → **DHCP** → **Networks**
2. Click the Action icon  
  
next to the network you want (this automatically selects it) and select **Edit** from the menu. The Network editor appears.
3. Click the **Discovery** tab.
4. Child networks inherit their discovery default settings from their parent networks. Click **Override** to change the **Enable Discovery** setting. (The **Discovery Member** setting remains unchanged.)
5. Deselect the **Enable Discovery** checkbox, and then save the configuration.

## Discovering Devices and Networks

To start discovery on connected networks, complete the following:

1. Ensure that your appliances are licensed for discovery.
2. Add the needed seed routers to each Probe appliance, as described in [Defining Seed Routers for Probe Members](#).
3. Add the necessary SNMPv1/v2 and SNMPv3 credentials at the Grid level and/or Member/Probe level. For information, see [Configuring SNMP1/v2 Credentials for Polling](#), [Configuring SNMPv3 Properties](#), [Defining Seed Routers for Probe Members](#), and [Defining Seed Routers for Probe Members](#).
4. If necessary, add CLI Credentials, including device admin username/password tuples and Enable passwords, at the Grid level and/or Member/Probe level, as described in [Configuring CLI Discovery Properties](#) and [Testing SNMP and CLI Credentials](#).
5. If necessary, enable the use of DHCP routers and servers as seeds to increase device discovery, as described in [Configuring SNMP1/v2 Credentials for Polling](#).

6. If you have extensive end host Ethernet segments connected to Ethernet switches, enable switch port discovery, as described in [Defining Seed Routers for Probe Members](#).

With these settings, the Probe appliances automatically begin discovering network infrastructure devices. You can elect to immediately discover or schedule discovery of new objects that you create and enable under IPAM or DHCP. Objects that allow immediate discovery include the following:

- **IPv4 Fixed Address** (see [Configuring IPv4 Fixed Addresses](#) for the complete procedure). You can **Enable Immediate Discovery** or **Exclude from Network Discovery** after creating the IPv4 fixed address, and override the SNMP credentials if necessary.
- **IPv6 Fixed Address** (see [Configuring IPv6 Fixed Addresses](#) for the complete procedure). You can **Enable Immediate Discovery** or **Exclude from Network Discovery** after creating the fixed address, and override the SNMP credentials if necessary.
- **IPv4 Reservation** (see [Configuring IPv4 Reservations](#) for the complete procedure). You can **Enable Immediate Discovery** or **Exclude from Network Discovery** after creating the IPv4 reservation, and override the SNMP credentials if necessary.
- **Host** (see [Adding Host Records](#) for the complete procedure). You can **Enable Immediate Discovery** or **Exclude from Network Discovery** after creating the host, and override the SNMP credentials if necessary.
- **IPv4 Network** (see [Configuring IPv4 Networks](#) for the complete procedure). You can **Enable Immediate Discovery** (option is enabled by default) and override inherited discovery **Polling Options** for the new network.
- **IPv6 Network** (see [Configuring IPv6 Networks](#) for the complete procedure). You can **Enable Immediate Discovery** (option is enabled by default) and override inherited discovery **Polling Options** for the new network.
- **IPv4 DHCP Range** (see [Configuring IPv4 Address Ranges](#) for the complete procedure). You can **Enable Immediate Discovery** (option is enabled by default) and override inherited discovery **Polling Options** for the new IPv4 DHCP range.
- **IPv6 DHCP Range** (see [Configuring IPv6 Address Ranges](#) for the complete procedure). You can **Enable Immediate Discovery** (option is enabled by default) and override inherited discovery **Polling Options** for the new IPv6 address range.

During configuration, you can choose to **Exclude from Network Discovery** if you wish to postpone discovery for specific object types.



#### Note

Individual IP addresses within a network, and specific object types (IPv4 reservation, fixed address, and host), may be excluded from discovery. You must explicitly select **Enable Discovery** for other object types (IPv4 and IPv6 Ranges, IPv4 and IPv6 Networks); you can optionally **Enable Immediate Discovery**. If you choose not to perform immediate discovery, but do **Enable Discovery**, the new network or other object is discovered at a normal time determined by Network Insight.

You can manually perform discovery on any object at any time by selecting the object and choosing

### Discover Now

**Discover Now** from the Toolbar. For more information, see [Object](#). When you do so, you see a status icon appear in the **Discover Now** data column

for the object under IPAM, in the **Data Management → DHCP** page and other locations.

By default, Grid discovery settings are the prevailing settings for all newly created objects. You can override basic discovery polling options for networks and DHCP ranges allowing immediate discovery.

In such cases, local settings take priority. Credentials cannot be overridden for networks and DHCP ranges,

## Using Discover Now to Discover an Existing Object



### Note

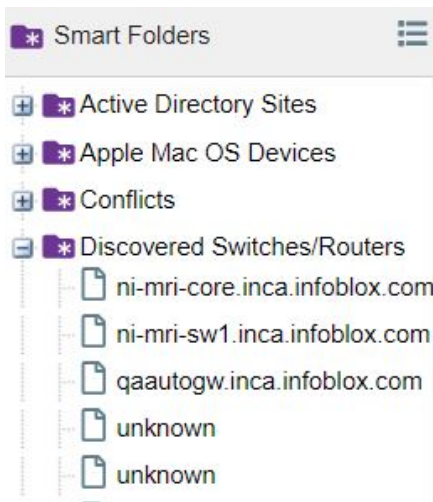
If after you select a network, object or IP and the **Discover Now** button is not enabled, make sure the network or other object has a discovery Probe member assigned to it.

After you create any supported IPAM object, you may wish to perform discovery on it at a later time. You can simply select the object and discover it.

1. From the **Data Management** tab, select the **IPAM** tab. The IPAM Home page appears.
2. Select the network or other object over which you want to perform discovery. Depending on the object type, navigate from the network level to the individual IP table in the **List** page to locate the object for immediate discovery.
3. Expand the Toolbar and click **Discover Now**. You can also click the Action icon

for the network and choose **Discover Now** from the menu.

The Probe member associated with the network or other object initiates a discovery procedure. Smart Folders and Discovered Devices



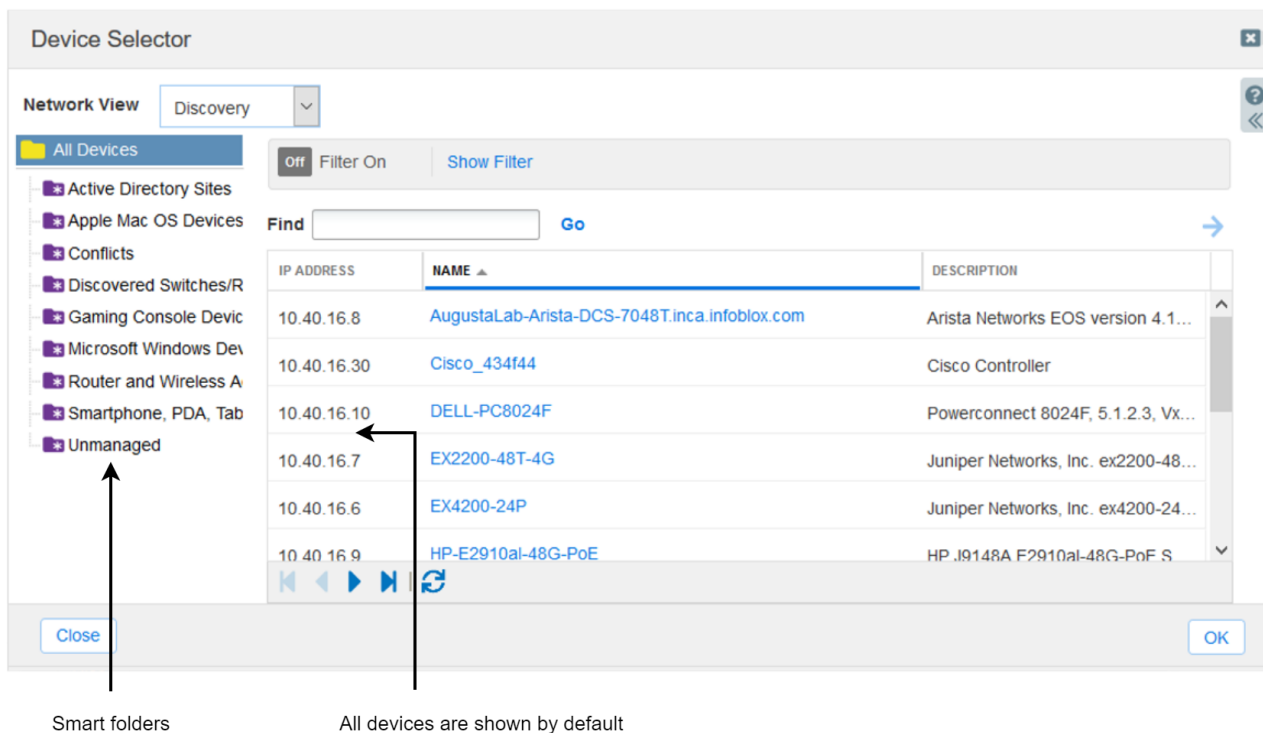
Grid Manager maintains a Smart Folder titled **Discovered Switches/Routers**, under which appears a list of all routers, switches and switch-routers that thus far have been discovered and catalogued through discovery. Clicking a device name opens the device's main page, with **Interfaces**, **Networks**, **IP Addresses**, **Assets** and **Components** panels. For more information, see [Managing Discovered Data](#).

Open the Smart Folders category under the Finder menu and click on the **Discovered Switches/Routers** folder. Clicking on a device name opens the device page under **Data Management** → **Devices** and shows the **Interfaces** page for the chosen device. For related information, go to [My Smart Folders and Predefined Smart Folders](#).

## Using the Device Selector

You use a dedicated Device Selector window to choose a discovered device for creating a port reservation with various IPAM objects such as IPv4 reservations, fixed addresses, host records, and provisioning of IPv4 and IPv6 networks on device interfaces; testing SNMP or CLI credentials, and other purposes, as shown in the below figure.

*Using the Device Selector Window*



Smart folders

All devices are shown by default

The chosen device is discovered and listed in the Devices panel in IPAM, or any other device on the network under the **All Devices** category in the left pane of the device selector.

Clicking a managed device's device name selects the device and brings you back to the originating page. Otherwise, select a device and click **OK** to continue.

If you have a long list of devices even after selecting a smart folder, enter a device name search value or a search expression in the **Find** field and click **Go**.

### Tips for Quick Navigation

To locate interface and device information quickly, along with associated IPAM objects that may be associated with elements such as port reservations, you can use Device and Interface terms in Global Search. (For additional information on Global Search, see [Global Search](#).)

Smart Folders provide another means of locating items such as IP addresses and IPAM objects of various types. Smart Folders provides additional filters for Device Type, Vendor, Model and Version. In a Smart Folder containing interfaces, you can filter by Admin Status, Operation Status, Trunk Status, VLAN ID or VLAN Name, and Description.

Using Smart Folders, you can also isolate all objects of a certain type by creating a smart folder with settings such as: **Type equals IPv4 Fixed Address**. Title the smart folder appropriately, to make clear what data set it is presenting.

### Starting Discovery

To ensure a successful discovery, complete the following:

1. In the Grid, install valid Discovery licenses on the Network Insight supported members that will later become the Consolidator and Probes. For information about the available license types and procedure to obtain and install licenses, see [Managing Licenses](#).
2. When you join the discovery members to the Grid, the first member automatically becomes the Consolidator while the others become Probes. If you want to change the roles of these members after they join the Grid, you can re-define their member types, as described in [Defining the Discovery Member Type](#). If you have only one

discovery member, it automatically becomes the Consolidator-Probe appliance after it joins the Grid. For more information about the Consolidator and Probes, see [Consolidator and Probes](#).

3. Configure applicable admin permissions for managing discovery and discovered data. For more information, see [Administrative Permissions for Discovery](#).
4. Define discovery interfaces and map them to corresponding network views. This step is especially important for discovering VRF virtual networks. For more information, see [Mapping Discovery Interfaces to Network Views](#).
5. Configure Grid discovery properties such as defining polling settings, configuring SNMP and CLI credentials, and configuring automatic VRF mapping. For more information, see [Configuring Discovery Properties](#). Note that you can override the Grid settings at the member and network levels, except for automatic VRF mapping which can be configured only at the Grid level.
6. Optionally, you can configure seed routers and map them to the corresponding network views. You can also use the default gateways for associated DHCP ranges and networks as seed routers for discovery by selecting the **Use DHCP Routers as Seed Routers** in the **General -> Advanced** tab of the *Grid Discovery Properties* editor. For more information, see [Defining Seed Routers for Probe Members](#) and [Defining Advanced Polling Settings for the Grid](#).  
Note that you must map each discovery interface, seed, and VRF to its respective network view in order to have a successful discovery for virtual routing instances.
7. Specify a network view, network container, or network to be discovered. For more information, see [Discovering Devices and Networks](#).
8. Optionally, define IP address exclusions when you want to exclude certain IPs from a discovery for various reasons. For more information, see [Excluding IP Addresses from Discovery](#).
9. Define blackout periods when you do not want the appliance to perform discovery. For information about how to configure blackout periods, see [Defining Blackout Periods](#).
10. Start the discovery service on the Consolidator and Probes to begin discovery, as described in [Starting and Stopping the Discovery Service](#). The Probe members continue to discover network devices within the defined networks.

## Network Insight Architecture

Discovery appliances that only detect devices and collect device data are called Probes, which are members of the Infoblox Grid, separately dedicated to the tasks of polling and discovery of networks and devices. A separate appliance, called the Consolidator, aggregates and organizes all collected device information from the Probes and synchronizes with the Grid Master. For more information, see [Consolidator and Probes](#).

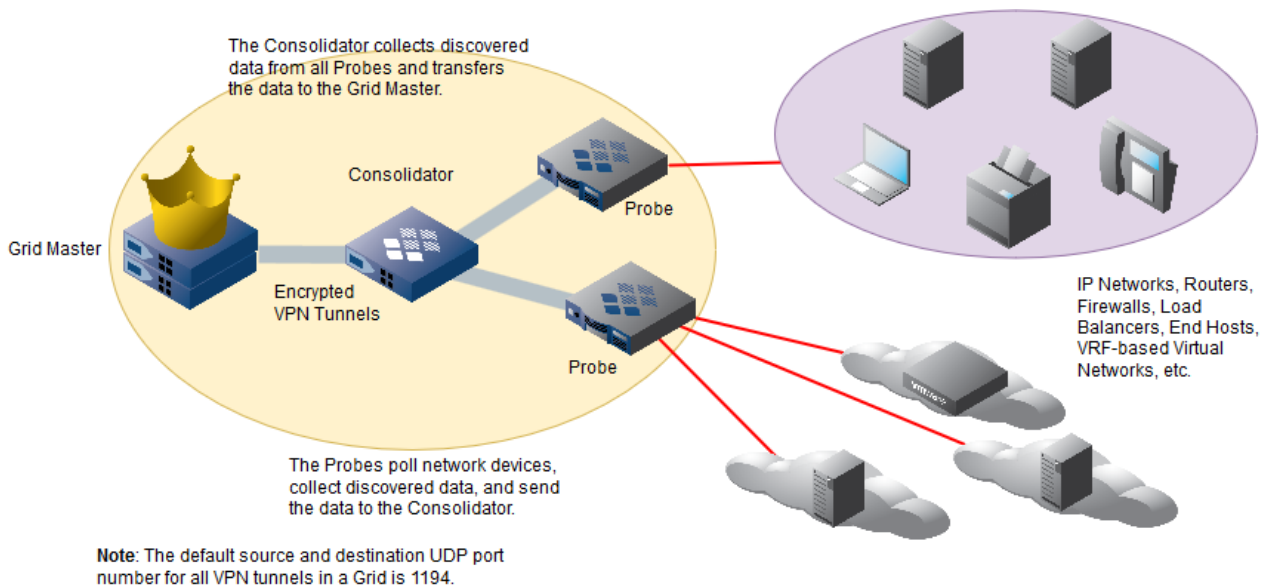
You can configure supported appliances as the Consolidator and Probes, as illustrated in the figure [Network Insight Appliances Added as Grid members](#) below. The Consolidator aggregates discovered data it collects from Probes and transfers the data to the Grid Master for device management and reporting purposes.

In the Grid, the Grid Master synchronizes data among all Grid members through encrypted VPN tunnels.

Communications between the Consolidator and Probes are also through encrypted VPN tunnels. The default source and destination UDP port number for VPN tunnels in a Grid is 1194. You can use the default port number or change it for VPN communications. Note that all the VPN tunnels in the Grid use the same port number you have chosen.

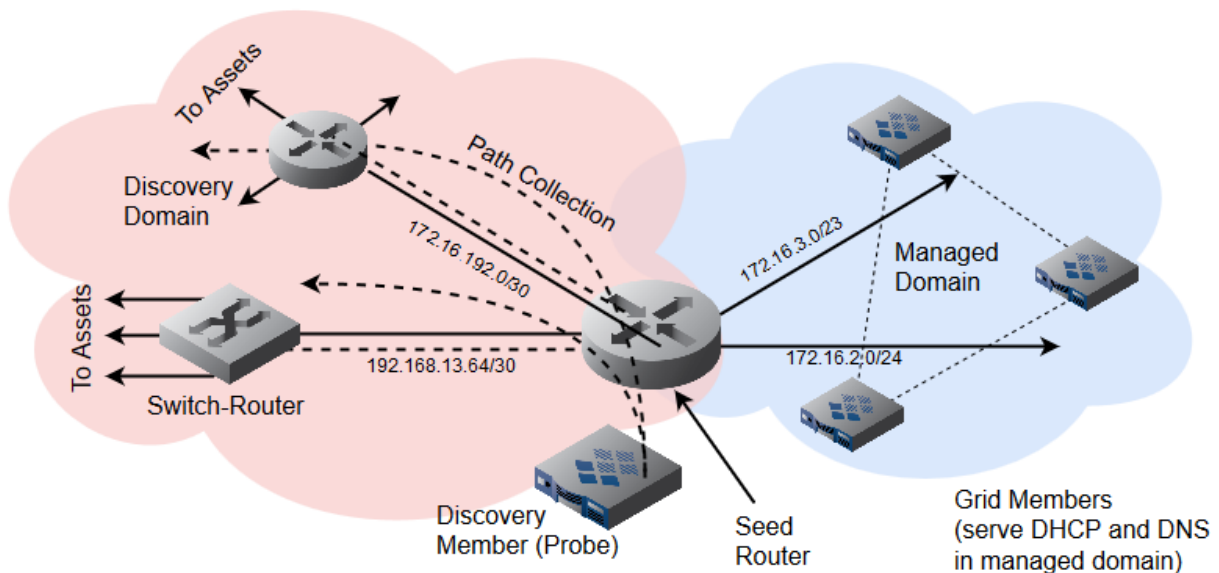
*Network Insight Appliances Added as Grid members*





Network Insight appliances use SNMP and other protocols to discover and catalogue a diverse assortment of device types, including the following: routers, enterprise switches, firewalls and security appliances, load balancers, enterprise printers, wireless access points, VoIP concentrators, application servers, VRF-based virtual networks, and end hosts. Network Insight provides a tool for administrators to gather key information about networks, including the discovery of routed paths and the host clouds behind enterprise switches, even in organizations where an Infoblox deployment already exists. In the figure Discovery in Action below, an appliance running discovery connects to an enterprise router, and uses its information to determine more about the networks that exist deeper within the unmanaged network, termed the discovery domain in this example.

*Discovery in Action*



As indicated in the above figure, discovery can trace through multiple hops and perform device discovery at every step, filling out the maps of unmanaged networks for the administrator. The collection of unmanaged network information extends to the networks of distribution Ethernet switches. Data collection also includes end hosts and application/file servers connected to edge switches in enterprise offices. Discovery uses the term assets to describe these devices. For more information, see [Viewing Assets Associated with Discovered Devices](#).

The Probes return discovery data to the Consolidator, which synchronizes device information with the Grid Master. Once



information about discovered networks and devices resides on the Grid Master, you can convert unmanaged networks and devices to managed objects, adding them to the NIOS database. For more information, see [Managing Discovered Data](#) and [About Automatic Conversion Rules](#).

You can also configure the appliance to send SNMP and email notifications when it discovers unmanaged devices and networks. For information about how to enable SNMP and email notifications for discovered unmanaged objects, see [Setting SNMP and Email Notifications](#). You can also manage these notifications by configuring the maximum number of unmanaged objects the appliance detects before it sends notifications and how often it notifies about these events. For information about how to configure these parameters, see [Defining Seed Routers for Probe Members](#).

You provide one or more routers as seed routers to act as the initial gateways for discovering other networks and their devices in the discovery domain (an example appears in the figure *Discovery in Action* above). You can also use DHCP routers (e.g., routers serving DHCP leases) as seed routers to aid in faster discovery.

When you create new networks, you can optionally provision them onto devices and perform discovery on them. Once you create the network, discovery can locate, poll and catalogue the network devices comprising the networks. This information is then synchronized with the Grid Master. For more information, see [Discovering Devices and Networks](#).



#### Note

For comprehensive coverage of port control features in Grid Manager, see [Pest Control Features in Network Insight](#) and its various subsections.

You can also exclude networks and IP addresses from discovery. The basic principle is that some devices do not need to be discovered, perhaps because they are already managed as part of a Grid and hence should not be subjected to discovery; because a device does not support SNMP; or for other organizational reasons. In the figure *Discovery in Action* above, networks 172.16.2.0/24 and 172.16.3.0/23 are excluded from discovery because they are already fully managed by a Grid. For more information, see [Excluding IP Addresses from Discovery](#).

You can define scheduled time periods in which Network Insight does not perform discovery operations in the network. These time periods are called *discovery blackouts*. All protocols associated with discovery (SNMP, CLI through Telnet and SSH, port scanning, fingerprinting and Ping sweeps) can be shut off during discovery blackout periods. This prevents discovery protocols from occupying network bandwidth during periods of peak usage. Network Insight does not communicate with devices in any way during a discovery blackout period. For more information on discovery blackouts, see [Defining Blackout Periods](#). Network Insight also provides a second type of blackout period for port configuration tasks, during which no tasks to change device port settings will execute. For more information, see [Defining Port Configuration Blackout Periods](#).

Related topic

[Infoblox Network Insight](#)

## Mapping Discovery Interfaces to Network Views

Discovery Probe members must have a designated interface or interfaces over which all discovery traffic exchanges take place. You may designate one or multiple discovery interfaces on each Probe member. You can specify the LAN1, LAN2, MGMT, and VLAN ports as the discovery interfaces. These ports must first be defined in the member network settings before they can be used for discovery. Note that when you configure VLANs on Network Insight appliances, the VLANs are used for discovery only. They do not support other core network services. On ND-1405, ND-2205, ND-4000, ND-V1405, and ND-V2205 Network Insight appliances, you can assign multiple VLAN interfaces on the same Probe to different network views.

To discover VRF virtual networks, you must discover their corresponding routes. In order to identify discovered data by routes, you associate each discovery interface with a network view. You must have a network view configured for each discovery interface, and you cannot share the same network view with multiple discovery interfaces on the same Probe member. However, discovery interfaces on different Probe members can share the same network view.

To map discovery interfaces to their respective network views, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Discovery**.
2. In the **Services** tab, select the checkbox of the Probe member you want to configure, and then Click **Edit** → **Member Discovery Properties** in the Toolbar.

3. In the *Member Discovery Properties* editor, select the **General** tab and complete the following:
  - **Discovery Interfaces:** Discovery members must have a designated interface or interfaces over which discovery traffic takes place. By default, appliances use the LAN1 port for discovery traffic. You may designate other ports such as the LAN2, MGMT, and VLAN ports as the discovery interfaces. These ports must first be defined in the member network settings before they can be used for discovery. You cannot modify the discovery interface settings while the discovery service is running on the appliance. Note that the VLAN interfaces you configured on any Network Insight appliances are used for discovery only. Other services are not supported on these appliances for VLANs. For information about how to define network interfaces on the appliance, see [Configuring Ethernet Ports](#). The **Discovery Interfaces** table displays all interfaces you have configured on the member. To discover using multiple interfaces, you must associate each interface with an available network view. A single default network view exists in NIOS by default. All networks created or discovered for NIOS management must be part of a network view. If more than one network view exists in the Grid, you can map a network view other than the default view to the interface. This essentially serves to allow one or more discovery members to perform discovery on separate routing domains, because a network view is comprised of a single routing domain with its own networks. If you do not want to use a configured interface as the discovery interface, simply leave the network view empty or unassigned for that interface. When you first set up an interface, no network view is assigned to the interface by default. The appliance displays the following for each interface you configure on the Probe member. To modify the network view for an interface, click the **Network View** column and select the network view you want to associate with the corresponding interface.
  - **Interface:** Displays the name of the interface. You cannot modify this field. Discovery supports the LAN1, LAN2, MGMT, and VLAN interfaces. You must first define these interfaces for the member before using them for discovery.
  - **VLAN Tag:** Displays the VLAN tag or ID for the corresponding VLAN interface. This field is left blank for all physical interfaces. You cannot modify this field.
  - **Network View:** Displays the current network view with which the interface is associated. An interface is not associated with any network view and this field is left blank by default, which means the interface is not used as a discovery interface. To modify the network view, click the **Network View** column for the corresponding interface and select the network view you want to reassign to the interface from the drop-down list. The appliance associates an interface with the default network view if you have not configured additional network views.
4. Save the configuration.

## Managing Discovery

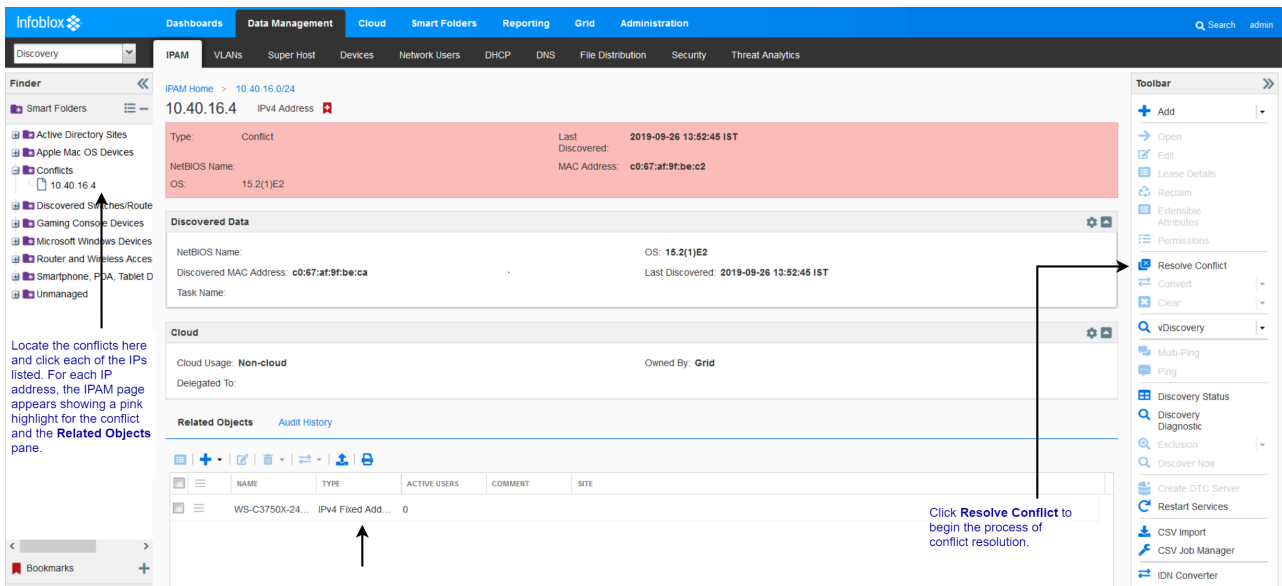
After you start the discovery service, you can do the following to manage the discovery and discovered data:

- Monitor the discovery status, as described in [Viewing Discovery Status](#).
- Execute discovery diagnostics to test the connection of a discovery member, as described in [Executing Discovery Diagnostics](#).
- Stop discovery on a specific network, as described in [Disabling Discovery for a Network](#).
- View a complete list of discovered devices, their associated interfaces, networks, IP addresses, assets, and components. For more information, see [Viewing Discovered Devices and their Properties](#).
- Resolve conflicts for discovered data, as described in [Conflict Resolution in Network Insight](#).
- Provision and de-provision networks and manage port configurations, as described in [Port Control Features in Network Insight](#).

## Conflict Resolution in Network Insight

You can sometimes encounter conflicts when defining port reservations for IPAM-managed objects such as Fixed IP addresses or host records. The quickest way to locate any conflicts in Grid Manager is to open the Conflicts Smart Folder as noted in the below figure.

*Locating conflicts and beginning their resolution*



Numerous types of conflicts are possible:

- Device Information conflict
- Port Reservation conflict, including Used Port Reservation conflicts (usually resulting from a request to reserve a port that has already been assigned to another IPAM object)
- Fixed address conflict
- IP Address conflicts
- DHCP Range conflicts (such as: *Discovered address is within an existing DHCP range but does not match an existing lease, fixed address, or exclusion range*)
- MAC Address conflict (such as: *Discovered MAC Address conflicts with existing fixed address*)



**Note**

When you execute discovery (**Discovery** -> **Discover Now** from the Toolbar), the appliance does not send SNMP trap if it finds any conflicting information between the NIOS data and the discovered data.

The Conflict Resolution wizard automatically recognizes the object associated with the conflict (which is listed in the **Related Objects** pane as noted in the figure Locating conflicts and beginning their resolution above) and ensures that changes you make during resolution are applied correctly to the object. An example appears in the below figure.

### Conflict resolution example

View/Resolve Conflict for 10.40.16.4

Description Discovered MAC address conflicts with existing fixed address

	Existing	Discovered
MAC Address	11:11:11:11:11:11	c0:67:af:9f:be:ca
NetBIOS Name		
OS	15.2(1)E2	15.2(1)E2
Last Discovered	2020-02-28 08:06:52 IST	2020-02-28 08:06:52 IST

Change the configured MAC address to be the same as the discovered MAC address

Keep fixed address and ignore this conflict for the next 1 day(s)

Cancel OK

### Resolving Port Reservation Conflicts

Sometimes, administrators may accidentally request a device port to be reserved for an IP address when the port is already reserved for another object, or try to apply a different port to an object that already has a port reservation. When these cases arise, Grid Manager reports a *conflict*.

To resolve port reservation conflicts:

1. Click the link provided in the **Conflicts** Smart Folder.
2. Expand the Toolbar and click **Resolve Conflict**, as shown in the figure Locating conflicts and beginning their resolution above.  
The Resolve Port Reservation Conflict dialog opens, showing the differences between the reserved and discovered information.
3. Choose from the following options:
  - **Change the reserved port to be the same as the discovered port.**
  - **Keep the configured port reservation and clear the conflict for the next 1 day(s).**  
Note that in the **Grid Discovery Properties** → **Advanced** tab, the **Ignore Conflict Duration** setting governs the default time duration to ignore (clear) certain types of conflicts that may occur when defining IPAM objects that are associated with discovered and managed devices, interfaces, or IP addresses. Increments can be defined in **Hours** or **Days**. For more information, see [Defining Seed Routers for Probe Members](#).
4. Click **OK** to save changes.  
Note that for other conflict examples, see Resolving Multiple Conflicts below.  
Another category of conflicts involves incorrectly defined device information for the object:
  - The reserved Device Type information provided is different from the discovered vendor and device type (Router vs. Switch, for example).
  - The defined Device Vendor information does not match with the discovered information.
  - A User Port Reservation conflict occurs when an unmanaged IP address attempts to use a port that is already reserved by an IPAM object on a different IP address.

You can choose from the following options:

- **Change configured information to discovered information.**
- **Keep the current device configuration and clear the conflict for the next 1 day(s).**

In virtually all cases, replacing the configured information with the discovered information successfully clears the conflict; click **OK** to commit changes or to temporarily clear the conflict.

## Resolving Multiple Conflicts

You can define objects for IP addresses, attempt to apply a port reservation, or incorrectly specify a value such as a MAC address or a vendor name, and accidentally cause multiple conflicts after creating the new object.

To resolve multiple conflicts for a particular IP address, use a Resolve Multiple Conflicts wizard:

1. To quickly locate any conflicts, open the Smart Folders panel and open the **Conflicts** list.
2. Click the IP address for any entry in the Conflicts list. The IPAM page opens for the selected IP address, with the top panel highlighted in pink to indicate the conflict.
3. Open the vertical toolbar and click **Resolve Conflict**.
4. If multiple issues are involved with the conflict entry, the Resolve Conflicts wizard lists each of them as shown in the below figure.

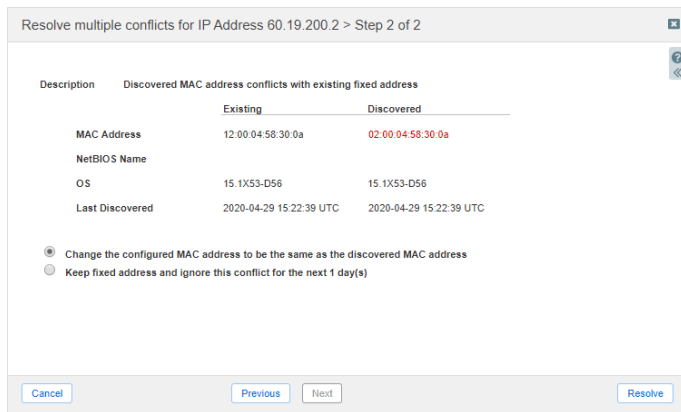
### *Resolution of Multiple Conflicts*



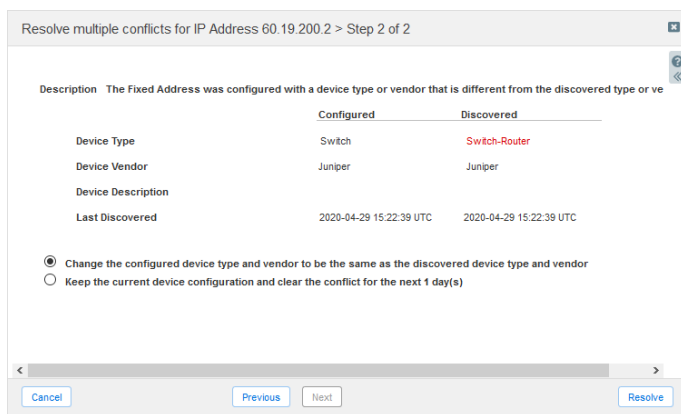
Multiple conflicts found. Select each in turn, in any order.

5. Select the conflict that you want to resolve first and click Next. For example, consider choosing to resolve the MAC Address conflict as shown in above figure. The second step of the wizard appears, listing the differences between the Existing and Discovered information for the conflict as shown in the figure below.

### *Troubleshooting the first selected conflict*



- a. In this case, the MAC address specified in the last fixed address object configuration, for that object, conflicts with the *discoveredMACAddress* associated with the IP. (You can verify this by checking the **RelatedObjects** tab in the IPAM page for the IP address.) Choose from one out of two options:
    - **Change the configured MAC address to be the same as the discovered MAC address;**
    - Keep fixed address and ignore this conflict for the next 1 day(s).
 In this example, the **Discovered** information for the MAC address associated with the Fixed Address object is one digit off from the **Existing** MAC information, which was entered incorrectly by the administrator. The Discovered MAC, shown in red, is the correct value and should be used to overwrite the record for the conflict.
  6. Select the **Update... with discovered data** option and click **Resolve**.
  7. The wizard updates with a return to the first step, in which you select the next conflict to resolve. A banner shows the result of the first resolution.
  8. Select the next conflict to resolve and click **Next**. As an example, the below figure shows a device configuration conflict.
- Device Information conflict for an object*



9. a. To resolve the conflict, the Configured information must be overwritten with the Discovered information. Choose any of the following:
  - **Change configured device type and vendor to be the same as the discovered device type and vendor**
  - **Keep current device configuration and clear the conflict for the next 1 day(s)**
 Other conflict types have similar options.
10. Select from the above choices and click **Resolve**.
11. Continue through the wizard to resolve the last conflict associated with the IP address.



#### Note

In the **Polling** tab → **Advanced** page of the *Grid Discovery Properties* editor, the **Ignore Conflict Duration** setting governs the default time duration to ignore (clear) certain types of conflicts that may occur when defining IPAM objects that are associated with discovered and managed devices, interfaces, or IP addresses. Increments can be defined in **Hours** or **Days**. This setting cannot be modified when resolving conflicts.

## Conflict Calculation Between IP Discovered vs DHCP Range and Leases

The Discovery Engine uses several background processes to determine the existence of IPs. Depending on the process the 'timestamp' of discovery (last-discovered date) corresponds to a real date by a direct check (ping, SNMP, CLI), or an indirect check (presence in forwarding tables, or ARP tables).

When the check is indirect the timestamp of an IP is NOT real time as compared to a direct check against the device. It means that this timestamp may not be accurate when the device was associated to this IP, but not really at the time denoted by the timestamp. This will have impact when computing conflict with RANGE and LEASE, as the timestamp may be inaccurately stated at a time the LEASE is already ended.

### Principle of Conflicts for DHCP Ranges

A DHCP Range is an object used for DHCP Service, to define which IP can be temporary allocated to a device using LEASE. It is expected that ONLY the device that hold the LEASE is using this IP.

If an IP is discovered on an IP slot of a DHCP Range and there is no corresponding LEASE, then there is a possible DHCP Range conflict (between the IP and the DHCP Range).

If in case the IP discovered is allocated for a LEASE, then we have two use case scenario:

- The LEASE IP has the same MAC than the discovered IP. All good.
- The LEASE IP has a different MAC than the discovered IP. Here a MAC ADDRESS conflict will be seen between the IP discovered and the IP LEASED.

### Computation of Conflicts (for RANGE or MAC ADDRESS Conflicts)

The conflicts are computed (raised or cleared) when:

- IP is discovered or re-discovered or has disappeared
- When the DHCP Range is updated (by UI, WAPI, PAPI)
- Any of the NIOS object that can change the slots in the RANGE are updated (Range reservation, Fixed Address, Reservation ...) by UI, WAPI or PAPI.



#### Note

LEASE operations DOES NOT trigger any recomputation of the conflict.

When computing a conflict, the process lookup for the IP discovered and verify this against any DHCP Range that includes the same IP:

- No Range : There will be no conflict with the Range (However, there can be other type of conflict see documentation)
- There is a Range, but there is nothing in that IP slot, neither Reservation, Fixed Addresses ... or LEASE: this IP is classified to be available for LEASE. If that IP is discovered, it means there is a device that uses that IP without holding the LEASE and a RANGE conflict is raised.
- There is a Range and this IP slot is used by:
  - Reservation or DHCP Reservation: no RANGE conflict.
  - Fixed Address or Host Address: no RANGE conflict. (There can be a MAC Conflict.)
  - LEASE:



- If the lease has currently no activity or never had any activity (eg BACKUP state). A RANGE Conflict is raised because we consider there is no LEASE.
- If the lease includes dates of activity (eg ACTIVE, RENEW, ABANDONED, FREE, ...), the “last-discovered” timestamp of the IP is checked against the dates of activity:
- If the “last-discovered” is BEFORE start of LEASE : we cannot say if there was a conflict or not. A SYSLOG “Maybe conflict - IGNORED” is issued, but no RANGE, neither MAC ADDRESS conflict is raised.
  - If the “last-discovered” is DURING activity of LEASE, it is expected. No RANGE conflict. However if the MAC does not match, a MAC ADDRESS conflict will be raised.
  - If the “last-discovered” is AFTER the activity of the LEASE : a RANGE conflict is raised, as there is no LEASE at time of the timestamp.



#### Note

Conflicts based on dates of “last-discovered” timestamp, depends on how the IP was discovered, with timestamps more or less accurate. Then a FALSE POSITIVE conflict can get raised. This condition needs to be manually verified along with the dates, assess the conflict and maybe resolve that conflict manually. If there is a delay to collect the LEASE from the DHCP Member, a conflict is raised at the time of processing the IP, because the corresponding LEASE was not yet consolidated on GM. This require a manual assessment of the conflict, and needs to be resolved manually.

## Executing Discovery Diagnostics

You can execute a discovery diagnostic to help determine why a specific device is presenting difficulties in discovery. For example, a given device may be reachable but show an overall status of **Failed** in the *Discovery Status* dialog. A discovery diagnostic steps through a complete discovery process based on the configuration on the Probe member to which the device is assigned. The diagnostic runs the gamut from fetching SNMP object ID information to ARP table reading and to ICMP pings and traceroutes.

You can do the following in the Discovery Diagnostics dialog:

- View all existing discovery diagnostic tasks that were executed in the last 12 hours.
- Enable or disable SNMP debugging for a device. The SNMP debugging is enabled by default.

Note that you must be a superuser to perform a discovery diagnostic. To execute a discovery diagnostic, complete the following:

1. From the **Data Management** tab, select the **IPAM** tab, and then click **Discovery Diagnostics** from the Toolbar.
2. In the *Discovery Diagnostic* editor, complete the following:
  - **Existing Discovery Diagnostic Task:** Select this option to choose an existing discovery task from the **Task ID** drop-down list. The appliance displays all the discovery diagnostic tasks that were executed in the last 12 hours.
  - **New Discovery Diagnostic Task:** Select this option to initiate a new discovery diagnostic. To start a new discovery diagnostic task, complete the following:
    - **Discovery Member:** Make sure that you select a discovery member for an IP address that does not exist in any network (in the **IPAM** tab) or excluded from discovery. Click **Select** to select a discovery member. You can click **Clear** to remove the discovery member
    - **IP address:** Enter the IPv4 or IPv6 address of the device on which you want to perform the test. The discovery diagnostic runs a full discovery procedure against the specified IP address. Ensure that you select the respective network view in which this IP address resides.
    - **NetworkView:** If you have multiple network views, select the network view in which the IP address resides from the drop-down list. If you have only one network view, which is the default view, the **NetworkView** drop-down list is hidden by default. NIOS conducts a discovery diagnostic for the IP address in the selected network view.
    - **CommunityString:** Specify the community string for the device if the required SNMP credential is not currently configured for the discovery member. It may not be necessary to enter a community string if the device is already discovered by NIOS and is a managed device.
    - **ForceTest:** To force a diagnostic against the device, select **Yes**.



- **Enable SNMP debug:** As a troubleshooting aid, the SNMP debugging option is enabled by default. When you enable this option, the appliance collects all SNMP communications between NetMRI and a device. The SNMP logs are useful for troubleshooting purposes. Clear this checkbox to disable the SNMP debugging for a device.
3. Click **Start** to start the discovery diagnostic. The lower pane displays the complete discovery sequence for the chosen device, and whether or not the discovery is successful. You can click **Stop** at any time to end the diagnostic sequence.

The output log of the diagnostic is displayed in the lower pane, and it shows the attempt for the complete discovery process. You can then do the following:

- You can click **Select All** to select the complete text in the lower pane for copying and pasting to the Clipboard and a text editor. You can also monitor the test messages in this pane.
- Click **Download as text** to download the complete discovery diagnostics in a text file for the selected device. The default name of the downloaded file is `discovery_diagnostic_nnn.nnn.nnn.nnn.txt`, where `nnn.nnn.nnn.nnn` is the IP address of the selected device. The **Download as text** button remains disabled until the download is complete.

## Viewing the Management State of IPs in Discovered Networks

You can view the management state for any IP address, in any network, that is associated with any discovered device.

1. From the **Data Management** tab, select the **Devices** tab. The Devices Home page displays a list of all devices currently found and catalogued by discovery.
2. Click the Action icon  
  
for a chosen device and choose **Networks** from the popup menu.
3. Choose a network from the list. Grid Manager switches to the IPAM page view of the selected network. The IPAM Home page displays the IP Map for the chosen network. The page shows information in graphical format, indicating elements such as **Used** Addresses, fixed addresses and IP reservations, **Unmanaged** IPs, **Host Not in DNS/DHCP**, and all other objects or information associated with IP management. The user benefits from this view by immediately seeing which IPs in the network contain devices that remain unmanaged by Grid Manager. These Unmanaged IP values appear in light yellow. Hovering the mouse over any IP address in the graphical table shows the information that has already been determined about the IP address.



### Note

An Unmanaged IP cannot be converted to Managed unless the network that contains it, is converted to managed status. For more information, see [Converting Unmanaged Networks under IPAM to Managed Status](#).

## Discovering VRF Virtual Networks

You can configure Network Insight to discover network devices that are configured or deployed within VRF (Virtual Routing and Forwarding) virtual networks. Using Network Insight to discover virtual networks provides visibility into your entire virtual network infrastructure, which allows you to view and manage overlapping IP addresses, VRF-specific data, and discovered end hosts. Note that a virtual network can consist of one or more physical devices that are configured to route packets using separate and distinct routing processes. Multiple routing tables can coexist on the same physical device or virtual device context, and traffic is exchanged among those devices using multiple routing tables. Depending on your network topology, there are a few ways you can use Network Insight for VRF network management. To use Network Insight effectively in the network, review the different deployment scenarios and configure Network Insight accordingly, as described in VRF Deployment Guidelines below. In addition, before you start a discovery for VRF virtual networks, ensure that you have reviewed the guidelines listed in *Special Considerations for Managing VRF Virtual Networks* below.

## Special Considerations for Managing VRF Virtual Networks

When you define discovery settings and perform management of VRF virtual networks, consider the following:

- If you limit the context of the SNMP community string in an individual VRF to the context of only that VRF, Network Insight will not be able to determine that the device it has discovered inside that VRF is the same device it has found inside other virtual networks. This may result in extra, un-correlated devices in the network. For information about how to configure SNMP credentials, see [Configuring SNMP1/v2 Credentials for Polling](#) and [Configuring SNMPv3 Properties](#).
- Network Insight will become aware of some devices inside of virtual networks from the route and ARP tables of routers that it manages. Without network connectivity into those virtual networks through a virtual discovery interface, Network Insight cannot discover all the devices or manage them. To create the necessary connectivity, you must configure a Network Insight discovery interface to be part of the VRF.
- Network Insight collects and parses the ARP and routing information from within a VRF context, but this data will not be used for further discovery unless the VRF virtual network is associated with a network view that is mapped on a discovery interface. For more information about how to map network views to discovery interfaces, see [Mapping Discovery Interfaces to Network Views](#).
- Global VRFs are labeled as: "default(IOS)" for IOS, "default" for Nexus and "master" for JunOS.
- For discovery and periodic polling on Juniper devices through an interface that is not in the Juniper default VRF (master), the query must use a special "default@credential" format. This setting assumes that users do not have management interfaces in a VRF. Your defined SNMP credentials for VRF-aware Juniper devices must use syntax similar to: "@credential". (Note that when querying VRF-aware Juniper devices via an interface that is in the default VRF, a plain community string can be used without the "@" character.)
- When configuring Network Insight to discover networks that employ route-leaking, discovery ranges for each network view should only be defined to include IP addresses that belong to that network view. In other words, any given Device IP should only fall within the discovery ranges of one network view. If discovery ranges are defined such that a device can be discovered by two different network views, the device may also be discovered via an unexpected network view. For information about how to define discovery ranges, see [Discovering Devices and Networks](#).

## VRF Deployment Guidelines

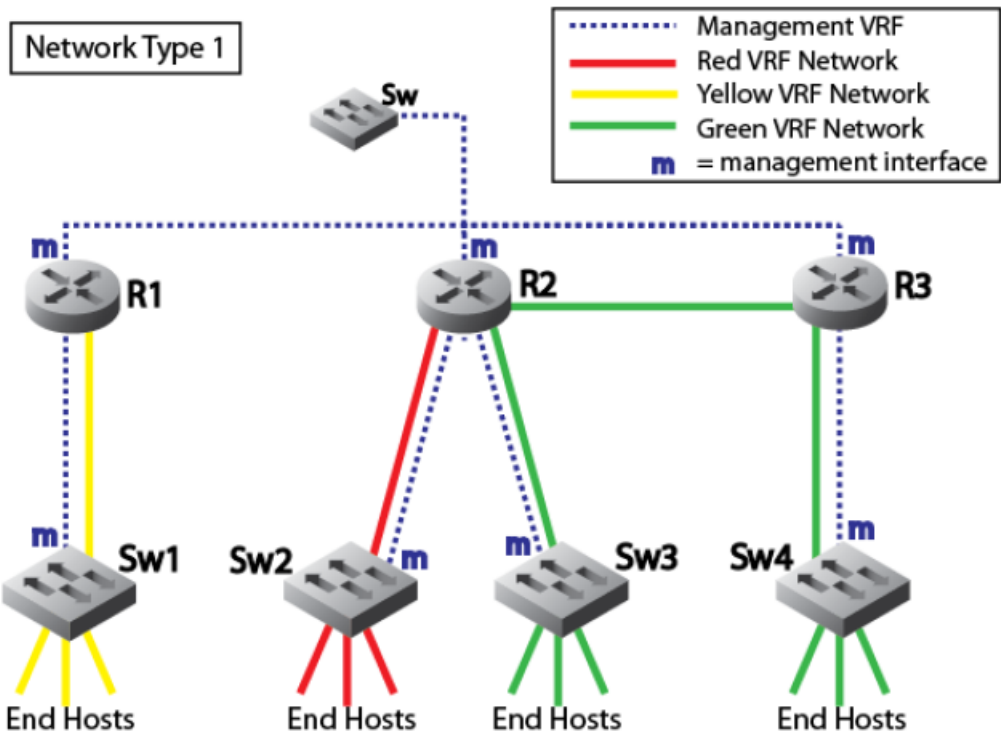
The topology of your network helps determine how you deploy Network Insight for VRF network management. To use Network Insight effectively in the network, you must possess some knowledge about your network so you can decide how to configure Network Insight to reach all the virtual networks you want to discover and manage. This section describes some common VRF-related network types for which you can deploy Network Insight.

Using the following three network types, all examples in this section help you define the number of network views and discovery interfaces so you can reach all locations in your network.

1. **VRF Network Type 1:** A network with a management VRF and several isolated production VRFs that include VRF-aware devices in the network.
2. **VRF Network Type 2:** A network with a shared service deployment VRF (shared VRF) and several isolated production VRFs that include VRF-aware devices in the network. The production VRFs share routes with the shared VRF, a practice also called route-leaking.
3. **VRF Network Type 3:** A network with several VRF-ignorant devices that reside in different L3 spaces, with no management VRF.

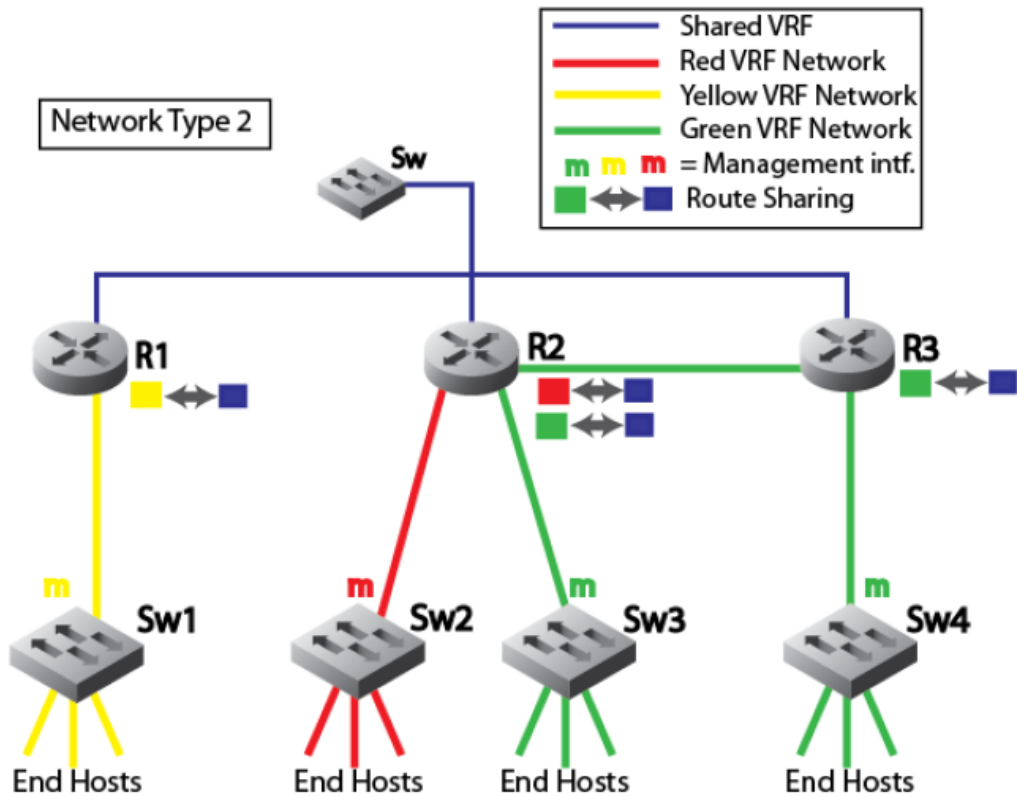
**VRF Network Type 1** has the following characteristics:

- A management VRF that reaches all VRF instances throughout the network.
- Isolated production VRFs (all VRFs can route to/from the management VRF but not to one another).
- The management VRF has complete visibility to all VRF instances in the network.




VRF Network Type 2 has the following characteristics:

- Uses a shared services deployment VRF to offer network services to the other production VRFs (shared VRF).
- All VRFs are reachable from the shared VRF, but VRFs cannot reach each other through the shared VRF or between each other.
- The production VRFs (Red, Yellow, Green) share routes with the shared services VRF (Blue).
- The shared VRF has complete visibility to all VRF instances.



VRF Network Type 3 has the following characteristics:

- Devices have management IPs only inside their respective networks.
- The routers in the network are VRF-aware; the switches are VRF-ignorant.



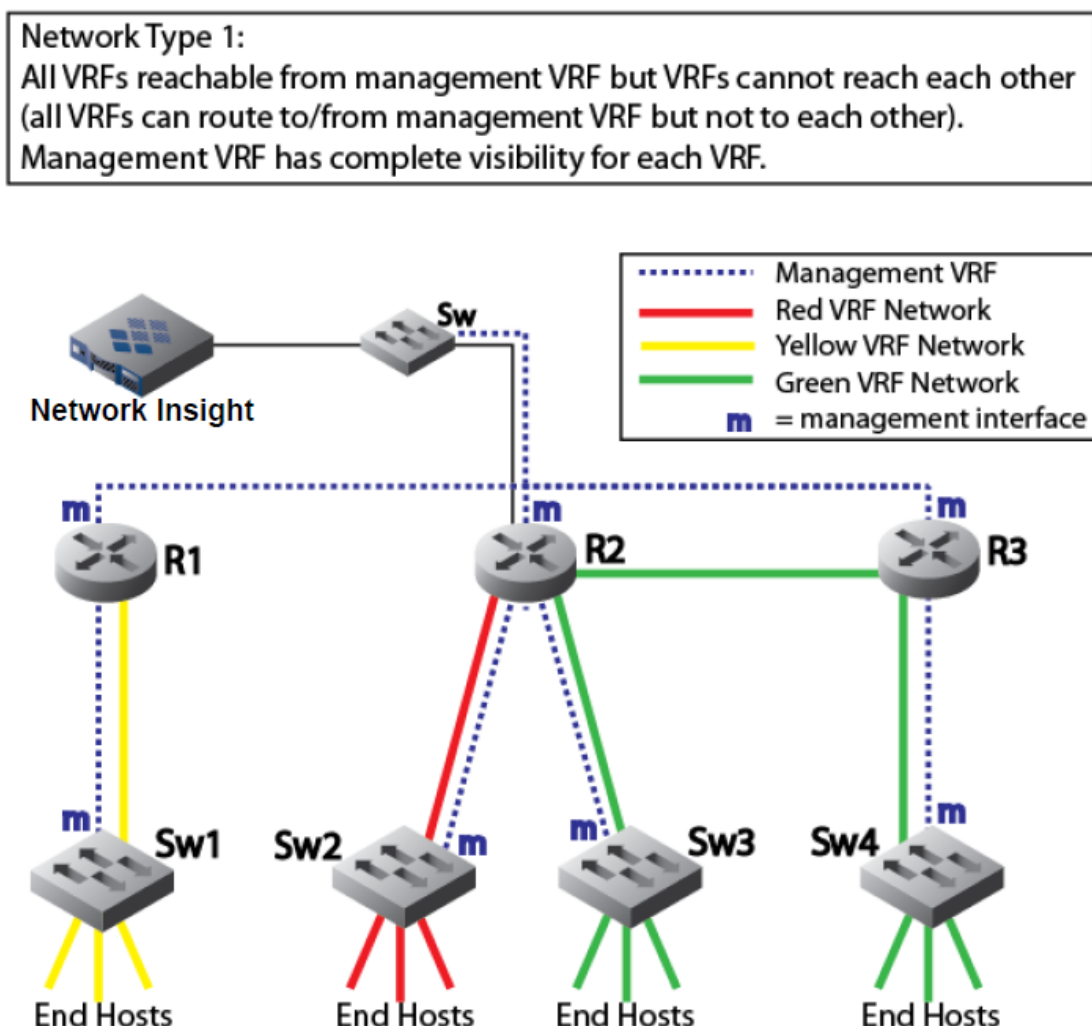
### Defining Network Views and Discovery Interfaces

In all three deployment types, you decide whether you want one or multiple network views based on how your network operates, as outlined in the three network types above. You can also consider the following guidelines:

- When all infrastructure devices for the network are reachable through a management VRF or a shared services VRF, and you do not need extended discovery capabilities to discover and/or manage end hosts, you can use a single network view. You also use a single virtual discovery interface to connect to the same 802.1q ID as the management VRF network. You can then discover and analyze all VRF-aware devices on the management VRF.
- If you want your devices end host and downstream device information separated for viewing and reporting, then you will want to use network views for each virtual network. Doing so is helpful for visual purposes, but it is not required.
- If you want ping sweeps, port scanning, fingerprinting and other discovery services into end hosts within each of your VRF networks, you must define multiple network views, one for each of your VRF networks; and each of which requires an associated virtual discovery interface and discovery ranges.

#### Deploying Network Insight in VRF Network Type 1: All Devices on a Management Network

The following figure shows an example of integrating Network Insight with **Network Type 1**. In this network deployment type, a single virtual discovery interface can manage all VRF instances' identification of ARP entries, because Network Insight needs only one discovery interface into the Management VRF.



#### Network Insight

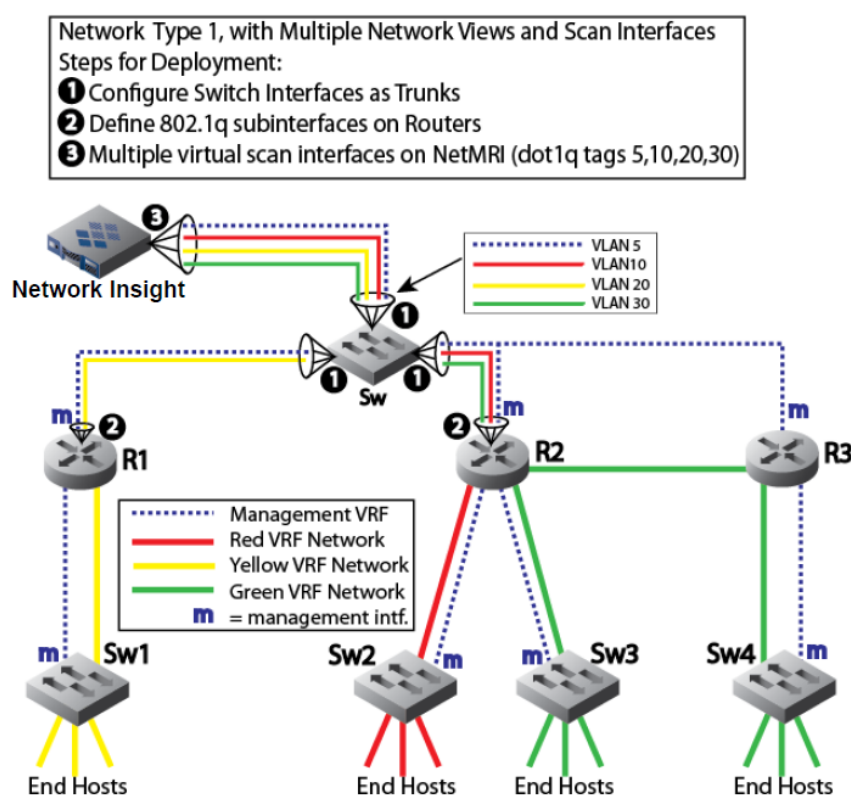
You configure the following for this example:

- **Network View:** Create one network view for the management VRF.
- **Discovery Interface:** Add the active discovery interface to the management VRF and tag it with the corresponding 802.1q VLAN value.
- **Discovery Ranges:** Define IP discovery ranges for the management network.
- All discovered VRFs must be associated with the network view configured for the management VRF.

### Deploying Network Insight in VRF Network Type 1: All Devices on a Management Network (Part 2)

The following figure shows the same topology for Network Type 1, but using multiple discovery interfaces and multiple network views.

In this example, the switch must be configured with the trunk port 'facing' Network Insight to forward Network Insight's tagged 802.1q traffic to the appropriate destination networks (VLAN 5, VLAN 10, VLAN 20 and VLAN 30 in this example). The encapsulated sub-interfaces are defined using the correct values on each port; the virtual discovery interfaces on Network Insight match these values.



### Network Insight

You configure the following for this example:

- **Network Views:** Create a network view for each network (Management, Red, Yellow, Green).
- **Discovery Interfaces:** Create virtual discovery interfaces for each VRF network.
- **Discovery Ranges:** Define IP discovery ranges for each VRF network.
- The discovered VRF instances must be associated with the network views to which they belong. For more information, see [Viewing Discovered VRFs and Mapping Network Views](#).

### Deploying Network Insight in VRF Network Type 2: All VRFs Reachable from a Shared Services VRF

This example illustrates the use of a shared service VRF between the distribution routers in the network and how Network Insight integrates into such a topology.

All virtual networks are reachable through a shared VRF, to which Network Insight may connect using a single virtual

discovery interface and reach all other VRFs from the one to which it is connected. Each Router in this topology also shares routes between the VRFs.

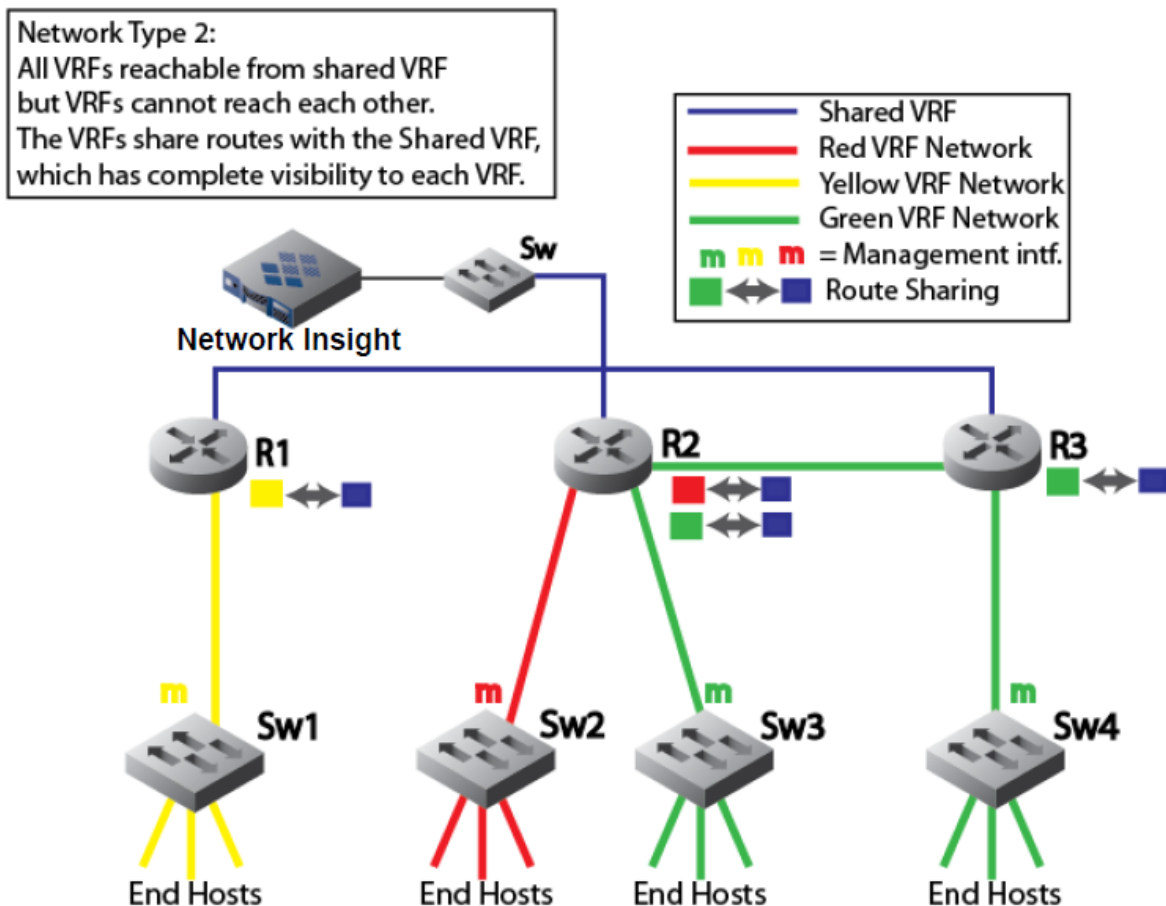
You configure the following for this example:

- **Network View:** Use one network view for the shared VRF.
- **Discovery Interface:** Create a virtual discovery interface on the network view. Use a single virtual discovery interface in Network Insight, and connect through the facing switch to the shared VRF using the tagged 802.1q value. There is a 1:1 ratio between network views and discovery interfaces.
- **Discovery Ranges:** Define IP discovery ranges in the single network view for all VRFs.
- All discovered VRFs must be associated with this network view.

If you want your device end hosts and downstream devices information separated, then use network views for each virtual network. This is helpful for viewing and reporting but it is not required. In this example, only a single network view is applied.

### Deploying Network Insight in VRF Network Type 2: All VRFs Reachable from a Shared Services VRF (Part 2)

In this version of the VRF Network Type 2 deployment, you use multiple network views and multiple discovery interfaces in a 1:1 ratio, with the same requirements for trunking and switch VLAN sub interfaces.



You configure the following for this example:

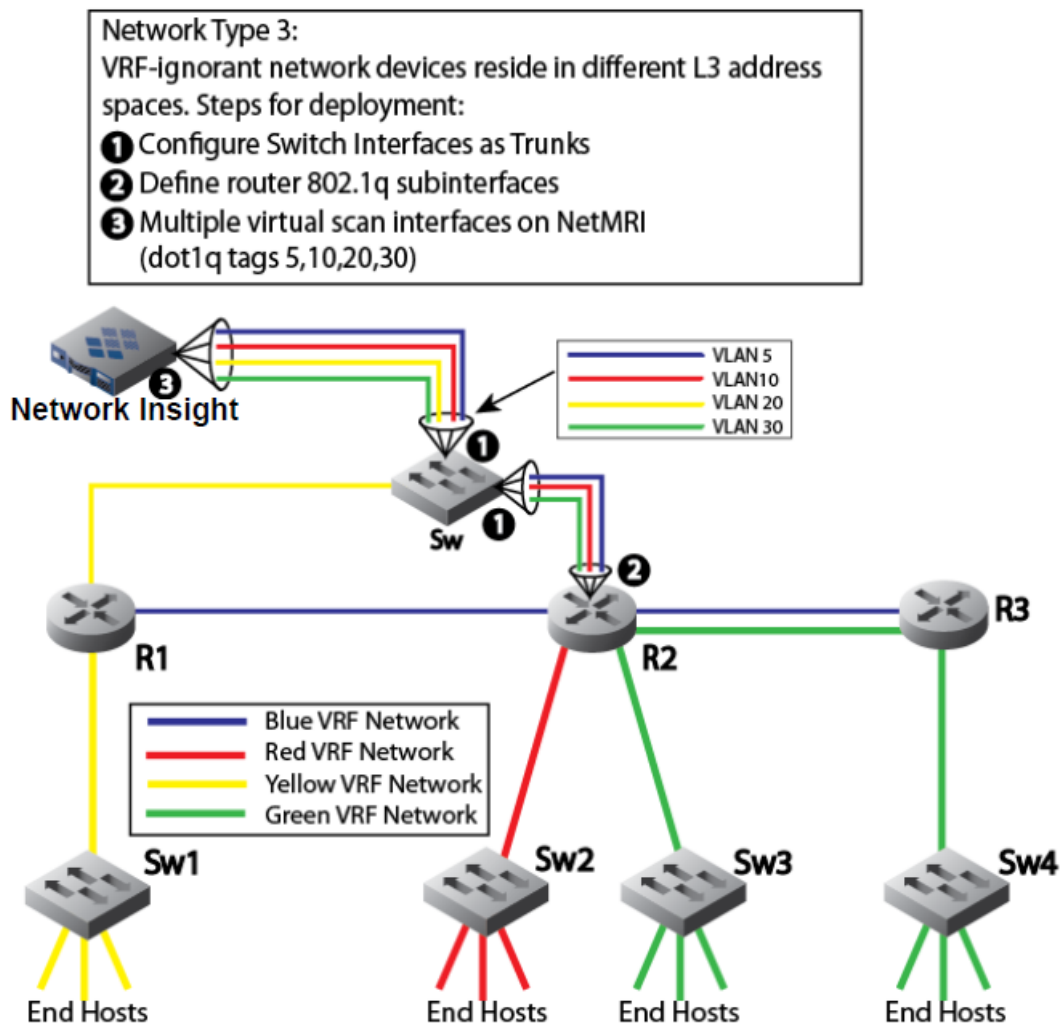
- **Network Views:** Create a network view for each network (e.g., Management, Red, Yellow, Green).
- **Discovery Interfaces:** Create virtual discovery interfaces for each VRF network.
- **Discovery Ranges:** Define IP discovery ranges in Network Insight for each VRF network.



- The discovered VRF instances must be associated with the network views to which they belong. For more information, see [Viewing Discovered VRFs and Mapping Network Views](#).

### Deploying Network Insight in VRF Network Type 3: Devices Reside in Disconnected Networks

In the final example, trunking is in use between Network Insight and its facing gateway switch into the managed network. This topology requires the use of multiple network views as all VRF networks are completely separate and cannot be reached through any management virtual network.



You configure the following for this example:

- **Network Views:** Create a network view for each network (e.g., Management, Red, Yellow, Green).
- **Discovery Interfaces:** Create virtual discovery interfaces for each VRF network.
- **Discovery Ranges:** Define IP discovery ranges for each VRF network.
- The discovered VRF instances must be associated with the network views to which they belong. For more information, see [Viewing Discovered VRFs and Mapping Network Views](#).

Each of the network views requires a single virtual discovery interface using 802.1q tagging as indicated in the figure. When defining the virtual discovery interfaces, use the 802.1q tag from the network devices. The primary differences are as follows:



- All devices do not have a management IP address in the so-called management VRF.
- The routers are VRF-aware while the switches are not.
- No VRF shares routes between any of the VRFs.

## Configuring Discovery Properties

To ensure a successful discovery, complete the following configurations for the Grid and Grid members that are acting as the Consolidator and Probes before you start a discovery:

- Define polling methods and schedule.
- Define advanced polling settings for TCP scanning and Ping sweeps. Also, specify routers and logging options.
- To collect data from SDN and SD-WAN devices, add and configure them as described in [Configuring Discovery for SDN and SD-WAN](#) below.
- If you use SNMP or CLI collection as the polling methods, define device credentials for data collection. .
- Assign credentials to device groups.
- Enable and schedule blackout periods for discovery and port configuration. For more information, see [Defining Blackout Periods](#).
- Configure automatic network view mapping for unassigned VRFs that have been discovered. For more information, see [Configuring Automatic VRF Mapping](#).
- Configure settings to monitor the lifecycle and vulnerabilities of discovered devices. For more information, see [Configuring Advisor Properties](#) below.

The following sections describe in detail how to configure discovery for the Grid, Grid members, and networks.



### Note

You must be a superuser to configure discovery properties for the Grid.

## Defining Basic Polling Settings

Grid polling settings apply to all Probe members and all discovery networks that are assigned to a Probe. You can override the Grid settings at the Probe member and network levels.

To define basic polling settings, complete the following:

1. For the Grid: From the **Grid** tab, select **Grid Manager > Discovery** service, and then select **Edit > Grid Discovery Properties** from the Toolbar.  
For Probe members: From the **Grid** tab, select **Grid Manager > Discovery** service, select a member, and then select **Edit > Member Discovery Properties** from the Toolbar.  
For networks: From the **IPAM** tab, click a network name, and then click the Edit icon.
2. In the *Grid Discovery Properties* or *Member Discovery Properties* editor, click **Polling > Basic**.  
In the *Network* editor, click **Discovery**.
3. If you want to override the inherited Grid settings for Probe members and networks, click **Override** and define the following settings.
4. **SNMP Collection**: Select this to execute SNMP protocols to discover and collect information such as traceroute/path collection, vendor and model, SNMP credential collection, routing and ARP tables, switch port data, and VLAN configuration data. If you disable SNMP collection, previously discovered data remains available for viewing. No new data is added and no existing data is removed.  
Some devices may not support SNMP, and some devices may not enable SNMP by default.  
Note When you disable SNMP collection on a network with enabled discovery, Network Insight still attempts to authenticate the SNMP credentials of devices that are newly discovered under this network. All newly discovered devices are automatically bound to a default group with enabled SNMP collection.
5. **CLI Collection**: Select this if you expect to use Network Insight to discover devices that support CLI connectivity through Telnet or SSH, and that you possess admin account information. NIOS can use device admin account logins to query network devices for discovery data collection, including IP configuration, port configuration, routing and forwarding tables, and much more.

Note that for SNMP and CLI Collection methods, configure device polling credentials in the **Credentials** tab of the editor. For more information, see Configuring Device Credentials below.

6. **Port Scanning:** Select this to probe the TCP ports. Ensure that you go to the **Advanced** tab to configure more settings for this option as described in the next section. If you disable Port Scanning, NIOS attempts no port probes other than SNMP on any device.
  - **Profile Device:** If enabled, NIOS attempts to identify the network device based on the response characteristics of its TCP stack, and uses this information to determine the device type. In the absence of SNMP access, the Profile Device function is usually the only way to identify non-network devices. If disabled, devices accessible via SNMP are identified correctly. All other devices are assigned a device type of Unknown. Profile Device is disabled by default for network polling.
7. **Smart IPv4 Subnet Ping Sweep:** Select this to execute Ping sweeps only on subnetworks that are known to exist but no IPs can be found within the subnet, such as through ARP or other means.
8. **Complete Ping Sweep:** Select this to enable brute-force subnet Ping sweeps on IPv4 networks. This method executes Nmap that uses ICMP echo requests, ICMP timestamp requests, and TCP SYN to ports 161, 162, 22, and 23 (for the SNMP, SNMPTRAP, SSH, and TELNET services correspondingly). Subnet ping sweeps are used as a last resort in the discovery process. Perform a subnet ping sweep if NIOS cannot identify any network devices in a given subnet. Subnet ping sweeps should be performed no more than once per day, and will stop on a given subnet once NIOS Discovery locates a network device and is able to collect data from it. Ensure that you configure advanced settings for this option in the **Advanced** tab as described in the next section.

Note that NIOS does not perform Smart Subnet Ping sweeps on subnets larger than /22. NIOS also does not perform Ping sweeps on IPv6 networks, because of the dramatically greater scale of network addresses in the IPv6 realm. The Complete Ping Sweep differs from the Smart Subnet Ping Sweep in the following ways:

  - The Complete Ping Sweep will run only against the specified range.
  - The sweep will run regardless of the range size.
  - The sweep will run regardless of the number of discovered devices within the specified range.
9. **NetBIOS Scanning:** Select this to enable NIOS to collect the NetBIOS name for endpoint devices in the network. This feature can be enabled only by users with SysAdmin privileges. This feature is globally disabled by default (and also for device groups) to prevent unexpected scanning of the network by a new collector.
10. **Automatic ARP Refresh Before Switch Port Polling:** Select this to enable refreshing of ARP caches on switches and switch-routers in the managed network before NIOS performs polling of switch ports. Enabling this feature applies only to switched Ethernet devices. This feature enables more accurate detection of all endpoint devices on L2 switches. Without ARP refresh, some endpoint devices may not be detected. This feature is globally disabled by default. Individual ARPs can also be set to enable or disable this feature.
11. **Switch Port Data Collection:** Select this to enable the Probe member to poll L2 enterprise switches. You can completely disable switch port polling by deselecting this checkbox. You can also separately schedule polling for switch port data collection as follows:
12. **Periodic Polling:** Define regular polling time periods. Choose a polling interval of 30 or more minutes or 1-24 hours.
13. **Scheduled Polling:** Schedule recurrent polling based on hourly, daily, weekly, or monthly time periods. Choosing this option, click the Calendar icon and a Polling Scheduler appears; click the Edit icon to make scheduling changes. Choose a recurrence pattern of **Once**, **Hourly**, **Daily**, **Weekly**, or **Monthly**. In all cases, you must choose an **Execution Time**.
14. Save the configuration.

## Defining Advanced Polling Settings

If you selected any polling settings that involve the TCP scanning and Ping sweeps, configure additional settings for them. Advanced SNMP polling settings consist of choosing the TCP Scan Technique, along with a number of specialized settings for Ping Sweeps and other operations.

To define advanced Grid-wide polling settings for TCP scanning and Ping sweeps, complete the following:

1. For the Grid: From the **Grid** tab -> **Grid Manager** tab -> **Discovery** service, select **Edit** -> **Grid Discovery Properties** from the Toolbar.  
For members: From the **Grid** tab -> **Grid Manager** tab -> **Discovery** service, select **Edit** -> **Member Discovery Properties** from the Toolbar.  
For networks: From the **IPAM** tab, select the *network* checkbox and click the Edit icon.
2. In the *Grid Discovery Properties*, *Member Discovery Properties*, or (*IPv4 or IPv6*) *Network* editor, click **Polling** -> **Advanced** and define the following settings.

3. If you want to override the inherited Grid settings for Probe members and networks, click **Override** and define the following settings.
4. **TCP Scan Technique:** Select the TCP technique you want to use for the discovery. The default is **SYN**. For more information, see [TCP](#).
  - **SYN:** Select this to quickly perform scans on thousands of TCP ports per system, never completing connections across any well-known port. SYN packets are sent and the poller waits for a response while continuing to scan other ports. A SYN/ACK response indicates the protocol port is listening while an RST indicates it is not listening. The SYN option presents less impact on the network.
  - **CONNECT:** Select this to scan IPv6 networks. Unlike the SYN option, complete connections are attempted on the scanned system and each successive TCP protocol port being scanned.
5. Specify the TCP ports settings:
  - In the table, select the checkboxes of the TCP ports you want to discover. To select all ports, click the checkbox in the header.
  - To add a new port, click the Add icon.
6. Specify other advanced polling settings:
  - **Purge expired assets data after:** Removes records of discovered assets that are no longer reachable after a specified period of time. The default is set to one day.
  - **Purge expired device data after:** Removes records of discovered network infrastructure devices that are no longer reachable after a specified period of time. The default is set to seven days, a more forgiving value given that devices sometimes require maintenance, upgrades or repairs, or in cases where hosts leave the network on long trips.
  - **ARP Aggregate Limit:** Sets a limit for the number of entries (IP addresses) per MAC address in ARP tables. If there are too many entries associated with a MAC address, this can be treated, for example, as a "honeypot". Therefore, MAC addresses with more entries than the specified limit are ignored and filtered out during data extraction and parsing. The default limit is 30 ARP table entries (IP addresses) per MAC address.
  - **Route Limit:** Limits the size of the routing table that discovery is required to collect from any given device. Some routers can have tables in the hundreds of thousands of entries, and collecting such a large body of data can impose performance problems in the network and in discovery data collection. This setting defaults to 3000, and automatically excludes BGP routes from the collection. Consult Infoblox Technical Support before making changes to this value.
  - **Ping Sweep Timeout (ms):** Period of time allowed, in milliseconds, before a Ping times out to any given device.
  - **Ping Sweep Attempts:** The number of attempts on each address in a Ping sweep before the sweep continues.
  - **Ping Sweep Frequency:** Defaults to 1, because ping sweep should not be executed more than once a day when the feature is enabled at the grid level or for a given discovery range. This setting affects the **Smart Ping Sweep** and **Complete Ping Sweep** features under *Grid Discovery Properties*.
  - **ARP Cache Refresh:** Defines the time period between ARP refreshes by Network Insight across all switch ports. Before any other switch port polling and discovery operations take place (including any global discovery polling operations initiated by the administrator), another ARP refresh is carried out by the Probe appliance regardless of the time interval. The default is five minutes, because switch forwarding tables are frequently purged from LAN switching devices. The default on Cisco switches is five minutes/300 seconds. Network Insight primarily uses ARP Cache refreshes to improve the accuracy of end-device discovery. Without this feature, some endpoints may not be discovered and cataloged.
  - **Ignore Conflict Duration:** Used when resolving conflicts and when choosing the option to Ignore the conflict when resolving it. The length of time during which conflicts is ignored is defined with this settings. Increments can be defined in **Hours** or **Days**.
  - **Number of discovered unmanaged IP addresses per notification:** The maximum number of unmanaged IP addresses that the appliance discovers before it sends SNMP and email notifications, if enabled. The appliance resets the counter after it hits this number and sends notifications. The default is 20.
  - **Interval between notifications for discovered unmanaged IP addresses:** This number determines how often the appliance sends SNMP and email notifications, if enabled, when it discovers the maximum number of unmanaged IP addresses (configured for **Number of discovered unmanaged IP addresses per notification** ). This is the time interval between two notifications for discovered unmanaged objects. Select the time unit from the drop-down menu. The default is five minutes.
  - **DNS Lookup Option:** Specify whether you want to perform a reverse DNS lookup from discovered IP addresses. Select one of the following from the drop-down list:

- **Network Devices:** Select this to resolve network device (switches and routers) IP addresses. This option is selected by default.
- **Network Devices and End Hosts:** Select this to resolve both network device (switches and routers) and end host IP addresses.
- **Off:** Select this to turn off reverse DNS lookups for discovered IP addresses.
- **DNS Lookup Throttle:** This is the value in a percentage that throttles the traffic on the DNS servers. Setting a lower value reduces the number of requests to DNS servers. You can specify a value between 1 and 100. The default value is 100.
- **Disable discovery for networks not in IPAM:** Disables executing discovery on any infrastructure networks that are not presented in the Infoblox IPAM system; e.g. present and managed in a network view or network container.
- **Authenticate and poll using SNMPv2c or later only:** For credential discovery and device polling exclusively using SNMPv2c and up, preventing use of SNMPv1, enable this checkbox.
- **Use DHCP Routers as Seed Routers:** Select this so the Probe members can use the default gateways for associated DHCP ranges and networks as seed routers to more quickly discover and catalog all devices (such as endpoint hosts, printers and other devices). All such default gateways are automatically leveraged by discovery, and no further configuration is necessary unless you wish to exclude a device from usage.  
Use this option carefully and avoid continuous updating of DHCP routers by a third-party component such as Microsoft servers, as it may trigger a discovery service restart when attempting to apply the new configuration.  
Ensure to check for a list of configured DHCP seed routers for any discovery Probe member in the **Seed** tab -> **Advanced** tab of the *Member Discovery Properties* editor.
- **Log IP Discovery events in Syslog:** Sends a message to the configured Syslog service when an IP address of an active host is discovered.
- **Log network discovery events in Syslog:** Sends a message to the configured Syslog service when a network discovery process takes place in the Grid.

7. Save the configuration.

## Configuring Device Credentials

Credentials apply to devices at the following levels:

- The Grid: Settings apply across the Grid and all Probe appliances licensed for Discovery.
- Discovery Probe appliances: You can use inherited Grid settings or override them.
- Individual devices: You can use inherited Grid or Probe settings or override them with device-specific settings.

You can configure the following types of device polling credentials:

- SNMPv1/v2 Credentials
- SNMPv3 Credentials
- CLI Credentials

Once configured, you can test the credentials.

For more information on configuring device credentials, see the following sections:

- [Configuring SNMPv1/v2 Credentials](#)
- [Configuring SNMPv3 Credentials](#)
- [Configuring CLI Credentials](#)
- [Defining CLI Credentials for Objects](#)
- [Testing SNMP and CLI Credentials](#)

If any SNMP or CLI credentials become obsolete, you can reset them for all affected devices at once. After that, Network Insight re-guesses the credentials for each device. This does not apply to CLI credentials manually set for specific devices. For more information, see the [reset snmp](#) and [reset cli](#) Administrative Shell commands.

You can assign a credential to a credential group that is specific to a particular device group. For more information about credential groups, see [Configuring Credential Groups](#) below.

## Configuring SNMPv1/v2 Credentials

An SNMPv1/v2 community string is similar to a password in that the discovered device accepts queries only from management systems that send the correct community string. This community string must exactly match the value that is entered in the managed system.

To add an SNMPv1/v2 credential, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Discovery**.
2. For the Grid: Click **Edit** -> **Grid Discovery Properties** in the Toolbar.  
For the Probe member: Select the *member* checkbox, and then click **Edit** -> **Member Discovery Properties** in the Toolbar.
3. Click the **Credentials** tab.
4. To override the inherited Grid settings for a Probe member, click **Override**.
5. Click the Add icon and specify the credential details in the corresponding cells:
  - **Read Community**: Enter a text string that the management system sends together with its queries to the network device during discovery.
  - **Credential Group**: For the Grid, select a group to which you want to assign the credential. For the Probe member, the table displays settings that were configured on the Grid, but only the default credential group is used for the member. You can edit the credentials list making up the default group by clicking **Override**.
  - **Order**: The order for attempting the use of the credential.
  - **Comment**: A text comment about the credential.
6. Optionally, you can test the credentials you added to the list. You can test SNMPv1/v2c and SNMPv3 credentials against any device or any IP address, at the Grid level or from any Probe member or network view. For more information, see *Testing SNMP and CLI Credentials* below.
7. Click **Save & Close** to save changes.

To find a specific string in the SNMPv1/v2 tab, enter the value in the **Go To** field and then click **Go**.

To remove a community string entry, select the checkbox and click the Delete icon.

To export the entire list of community strings in a table file readable by a spreadsheet program, click the Export icon and choose Export Data in Infoblox CSV Import Format. To export all data in a different format, click the **Export** icon and choose **Export Visible Data**.

## Configuring SNMPv3 Credentials

SNMPv3 allows the use of two secret keys for every credential — one for authentication and another for encryption. Network Insight allows flexible application of keys — authentication but no encryption, for example. You define users in one of the three following ways:

- SNMPv3 user with no authentication or privacy credentials.
- SNMPv3 user with authentication but no privacy credentials.
- SNMPv3 user with both authentication and privacy credentials.

To add an SNMPv3 credential, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Discovery**.
2. For the Grid: Click **Edit** > **Grid Discovery Properties** in the Toolbar.  
For the Probe member: Select the *member* checkbox, and then click **Edit** > **Member Discovery Properties** in the Toolbar.
3. Click the **Credentials** tab > **SNMPv3** tab.
4. To override the inherited Grid settings for a Probe member, click **Override**.
5. Click the Add icon and specify the credential details in the corresponding cells:
  - **Name**: The username for the credential.
  - **Auth Protocol**: Select one of the listed authentication protocols.
  - **Auth Password**: The password to use with the authentication protocol.
  - **Privacy Protocol**: Select one of the listed privacy protocols.
  - **Privacy Password**: The password to use with the privacy protocol.

- **Credential Group:** For the Grid, select a group to which you want to assign the credential. For the Probe member, the table displays settings that were configured on the Grid, but only the default credential group is used for the member. You can edit the credentials list making up the default group by clicking **Override**.
  - **Order:** The order for attempting the use of the credential.
  - **Comment:** A text comment about the credential.
6. Optionally, you can test the credentials you added to the list. You can test SNMPv3 credentials against any device or any IP address, at the Grid level or from any Probe member or network view. For more information, see [Testing SNMP and CLI Credentials](#) below.
  7. Click **Save & Close** to save changes.

To find a specific SNMPv3 entry, enter the value in the **Go To** field and click **Go**.

To remove an SNMPv3 authentication entry, select the checkbox and click the Delete icon.

To export the entire list of community strings in a table file readable by a spreadsheet program, click the Export icon and choose Export Data in Infoblox CSV Import Format. To export just the subset of data that is visible in the dialog, click the **Export** icon and choose **Export Visible Data**. A **Show Passwords** option allows the secret keys to be visible in the import.

### Configuring CLI Credentials

SNMP protocols provide a powerful means of querying devices for broad arrays of information. The CLI discovery feature is required for port control tasks including port configuration and network provisioning and de-provisioning, but is not used for other discovery operations or to otherwise manage devices. By default, Probe appliances inherit their member discovery properties, including CLI credential sets, from the Grid level. Enable passwords are entered in separate records and kept as a separate list in Grid Manager.

For CLI credentials, you define a global set of Admin account/password tuples, as well as Enable passwords, at the Grid level. You can also specify credentials and Enable passwords for individual devices at the member level. Should such a credential not work for a given device, or if command-line access is lost for a device, Network Insight re-guesses credentials from the Grid-level credential list, including vendor defaults if available.



#### Note

You can test username/password credentials or an Enable password credential. You can also combine a username/password credential and an Enable password credential as part of the same test.

To add a CLI credential, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Discovery**.
2. For the Grid: Click **Edit > Grid Discovery Properties** in the Toolbar.  
For the Probe member: Select a *member* checkbox, and then click **Edit > Member Discovery Properties** in the Toolbar.
3. Click the **Credentials** tab > **CLI** tab.
4. To override the inherited Grid settings for a Probe member, click **Override**.
5. Click the Add icon to add a new CLI username/password entry to the list. Select the **Credential Type**, which can be one of two choices.
6. In **Login Credentials**, click the Add icon and specify the credential details in the corresponding cells:
  - **Protocol:** Select **SSH** or **Telnet**. Infoblox recommends the use of SSH.
    - **SSH:** SSH credentials require both a username and a password. The default protocol is SSH.
    - **Telnet:** In Network Insight, Telnet credentials must use both a username and a password. Note that should you choose to use a Telnet-based credential, Network Insight requires both the username and password for the login account. This also applies when you override the CLI credentials on objects such as a fixed address, host, or IPv4 reservation. For more information, see the section [Defining CLI Credentials](#) *Settings for Objects* below.
  - **Name:** Username for the CLI login account.
  - **Password:** Login password for the CLI login account.
  - **Credential Group:** For the Grid, select a group to which you want to assign the credential. For the Probe member, the table displays settings that were configured on the Grid, but only the default credential group is used for the member. You can edit the credentials list making up the default group by clicking **Override**.
  - **Comment:** A text comment describing the CLI login account.



- **Order:** By default, Network Insight inserts the new credential record at the bottom of the credentials list, which is reflected by its **Order** value, showing the order used for attempting the use of CLI credentials. Enter a new value in the **Order** field if you want the new credential to be in a position other than the last in order.
7. In **Enable Credentials**, click the Add icon and specify the credential details in the corresponding cells:
    - **Protocol:** SSH or Telnet. Infoblox recommends the use of SSH.
    - **Password:** Enable password for device configuration access.
    - **Credential Group:** For the Grid, select a group to which you want to assign the current credential. For the Probe member, this setting is inherited and cannot be changed.
    - **Comment:** A text comment about the credential.
    - **Order:** By default, Network Insight inserts the new record at the bottom of the list, reflected by its **Order** value, showing the order used for attempting use of the CLI credentials. Enter a new value in the **Order** field if you want the new credential to be in a position other than the last in order.
  8. Optionally, you can test the credentials you added to the list. For more information, see [Testing SNMP and CLI Credentials](#) below.
  9. Click **Save & Close**.

### Defining CLI Credentials for Objects

You can define CLI credentials and enable password credentials for individual devices through associated IPAM objects:

- Fixed addresses
- IP reservations
- Hosts

For a quick way to locate all objects of a certain type in the Grid, you may create a smart folder with settings such as **Type > Equals > IPv4 Fixed Address**. Title the smart folder appropriately.

To define and test the CLI credentials for an IPAM object, complete the following:

1. From the **Data Management** tab, select the **IPAM** tab.
2. In the *IPAM IP List* page or the *IPAM IP Map* page, navigate to the required network and then to the IP associated with the object you want to edit.
 

Note for each network, the IP list page provides a **Type** data column showing the IPAM object type that is associated with any IP address. Also, check the **MAC Address** column in the IP List page for information about associated objects.
3. Click the IP address. On the IP address page, click the **Related Objects** tab.
4. Select the checkbox for the object in the Related Objects panel and click **Edit**.
5. In the object editor, click the **Discovery** tab.
6. Click **Override CLI Credentials**.
 

By default, CLI credential definitions use SSH at the object level. Select **Allow Telnet** if you want to allow both SSH and Telnet credential usage. Infoblox recommends SSH because of better security.
7. Enter the **Name** and **Password** values and the **Enable Password** value.
8. Click **Test CLI Credentials** to test the CLI discovery credential settings applied to the object.
9. When finished, click **Save & Close**.

### Testing SNMP and CLI Credentials

After configuring SNMP and CLI credentials, you can click **Test Credentials** in the SNMP Credentials or CLI Credentials panel to test the credentials. Credential testing ensures that the configured credentials work for as many devices and networks as possible. The procedure in this section applies to both the Grid and the member levels. You can override the Grid settings at the member level.

For CLI credentials, you can test an admin login name and password tuple as well as a following enable credential, if necessary. You can also override CLI credentials and enable credentials for IPAM objects such as fixed addresses, IP reservations, and host objects. You can test any credential set, an enable credential, or both in combination against any device within any network view. Network Insight sets the login sequence to match the command-line standards for the selected device.

To test SNMP credentials or CLI credentials, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Discovery**.
2. For the Grid: Click **Edit > Grid Discovery Properties** in the Toolbar.  
For a Probe member: Select a *member* checkbox, and then click **Edit > Member Discovery Properties** in the Toolbar.
3. In the **Grid Discovery Properties** or **Member Discovery Properties Editor**, click the **Credentials** tab > **SNMPv1/v2** tab, **SNMPv3** tab, or **CLI** tab.
4. Select the checkbox or checkboxes for the Login credentials and/or Enable credentials (if applicable) you want to test.  
For a Probe member, click **Override**, and then select the credential checkboxes.
5. Click **Test Credentials**.
6. In the **Test SNMP Credentials** or **Test CLI Credentials** dialog box, complete the following:
  - **IP Address:** Select this to test the credential against an IP address of a reachable device in a network (even if it has not been discovered or managed through NIOS) within a specified network view. Enter the IP address in the field. Ensure that you select the respective network view in which this IP address resides.
  - **Network View:** If you have multiple network views, select the network view in which the IP address resides from the drop-down list. If you have only one network view, which is the default view, the **Network View** drop-down list is hidden by default. NIOS conducts credential testing for the IP address in the selected network view.
  - **Device:** Select this to test against a selected device. Click **Select Device** and the *Device Selector* appears. If you have multiple network views, you must first select the network view in which the device resides from the **Network View** drop-down list, located at the upper left-hand corner of the selector. If you have only one network view, which is the default view, the **Network View** drop-down list is hidden by default. You can check the device categories under **All Devices** to locate discovered switches and routers or any unmanaged devices previously detected by discovery under NIOS. You can explore categories including Discovered Switches/Routers, Microsoft Windows Devices (this can include items such as Windows Servers of various types), Router and Wireless Access Point Devices, Unmanaged, and others. By default, all devices previously discovered appear in this selector. If you have a long list of devices, you can enter a device name search value or a search expression in the **Find** field and click **Go**. You can also click **Show Filter** to narrow down your selection by selecting available filters. Click **OK** after you have selected a device and its corresponding network view.
7. Click **Start** to begin testing the credential against the IP address or selected device. The communication and testing processes appear in the lower panel of the editor.



#### Note

If the specified IP address is excluded from all discovery ranges, if it is not a part of the selected network view, or if the credential is entered with missing information, a message appears at the top of the editor after clicking **Start**. Otherwise, the test begins and its process and results appear in the lower panel of the editor.

## Configuring Credential Groups

You can organize specific credential sets into groups that you can use to guess device credentials by assigning different credential groups to IPAM objects such as networks. For example, when a credential group is assigned to a network, Network Insight runs the guessing process only over credentials of this group for devices in this network.

You can use credential groups for credential guessing on the following levels:

- The Grid
- Probe members
- IPAM networks and network containers
- DHCP networks and DHCP ranges

A default credential group preexists on the Grid. The default credential group is automatically pre-assigned to the Grid, members, and IPAM objects such as networks, network containers, or ranges. You can create custom credential groups and add credentials to them. After that, you assign credential groups to the Grid, Probe members, and IPAM objects.

For more information, see the following sections:

- [Creating a Credential Group](#)



- Adding SNMP and CLI Device Credentials to a Credential Group
- Assigning a Credential Group
- Viewing Credential Group Used for Device

### Creating a Credential Group

You can create, edit and delete custom credential groups. As for the default credential group, you can rename it, but not delete it.

To create a credential group, complete the following:

1. Go to **Grid > Grid Manager > Discovery**.
2. In the **Toolbar**, click **Edit > Grid Discovery Properties**.
3. Select the **Credential Groups** tab.
4. Click the Add icon.
5. Type a name for the credential group and press **Enter**.  
Note that the **Name** field does not support the Unicode encoding.
6. Click **Save & Close**.

To edit a credential group, click the group name in the table cell and edit as required.

To delete a credential group, select the group and click the Delete icon.

### Adding SNMP and CLI Device Credentials to a Credential Group

You populate a credential group with credentials in the Grid Discovery Properties. For Probe members, you can edit the credentials list inside a group for the default credential group only. choose to inherit the Grid credential groups settings or to override them with the default group. If the default group contains credentials from both Grid and member and it is assigned to a network, Network Insight uses only member-level credentials.

See the following sections on how to add device credentials of different types to a credential group:

- [Configuring SNMPv1/v2 Credentials](#)
- [Configuring SNMPv3 Credentials](#)
- [Configuring CLI Credentials](#)

### Assigning a Credential Group

Probe members, networks and ranges inherit the credential groups assignment from the Grid. You can override this assignment with another credential group for networks and ranges. For members, you cannot assign a credential group as they always use the default group for credential guessing.

#### Assigning a Credential Group to the Grid

To assign credentials to the Grid, complete the following:

1. Go to **Grid > Grid Manager > Discovery**.
2. In the **Toolbar**, click **Edit > Grid Discovery Properties**.
3. On the **Polling** tab.
4. For **Credential Group**, select the required credential group.
5. Save the configuration.

#### Assigning a Credential Group to Network or Network Container

To assign credentials to an IPv4/IPv6 network or network container, complete the following:

1. Go to **Data Management > IPAM** or **DHCP**.
2. Click the name of an existing network or network container.
3. Click the Edit icon.
4. Click the **Discovery** tab.

5. **Credential Group:** The field displays the selection inherited from the Grid unless you override it. Click **Override** and select the required credential group.
6. Save the configuration.

Assigning a Credential Group to DHCP Network or Range

To assign credentials to a DHCP range, complete the following:

1. Go to **Data Management > DHCP**.
2. Click the name of an existing IPv4/IPv6 network.
3. Click the Add or Edit icon > **Range**.
4. Click the **Discovery** tab.
5. **Credential Group:** The field displays the selection inherited from the Grid unless you override it. Click **Override** and select the required credential group.
6. Save the configuration.

Viewing Credential Group Used for Device

To view which credential group was used to guess credentials for a particular device, complete the following:

1. Go to **Data Management > IPAM** or **Devices** tab > **Discovery Status**.
2. Hover the mouse over SNMP or CLI credential status of a device.  
A tooltip mentioning the credential group name appears.

Defining Seed Routers for Probe Members

Seed routers can be defined only on Probe appliances. You can define seed routers that NIOS uses in quickly performing network discovery. The definition of seed routers is highly recommended for IPv4 networks and is required for IPv6 networks. For the discovery of any IPv6 networks, you must use seed router values that comprised of at least one well-connected IPv6 router, preferably with routes to all other networks to be managed. In some cases, seed routers may not have the full routing tables or be unable to provide full information for some reason. The general rule of thumb is that more seed routers are better, but the connectivity of seed router(s) also helps determine how many seed routers you need. Avoid having more seed entries than necessary.

You must associate each seed router with a network view so the appliance can properly discover virtual networks when using multiple seed routers.



#### Note

All NIOS Probe members automatically use their default gateway as a seed router. These gateways are automatically displayed in the table. For effective use of seed routers, you must also provide SNMP credentials to NIOS to allow it to pull the key routing and connectivity information, including the IPv6 routing table and the local Neighbor Discovery Cache, from the device. If you do not define a seed router, it is recommended that you enable discovery for a network or DHCP range.

You can check **Discovery Status** to see whether a seed router is successfully being reached and whether the seed is providing information. By reviewing discovery status for each seed router, you can determine whether Network Insight should be able to discover the network successfully, or if there are possible configuration errors preventing network discovery, without having to wait to see what Network Insight finds. For seed routers, **Reached Status** and **Overall Status** should both read as **Passed**.

To add, view, or delete seed routers for a Probe, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and click **Discovery**.
2. Select the checkbox for any Probe appliance on the Discovery page and click **Edit > Member Discovery Properties** from the **Toolbar**.
3. In the **Member Discovery Properties Editor** dialog box, click **Seed**.
4. To add a seed router, do the following:

- a. Click the **Add** icon. The Grid Manager adds a row to the table.
- b. In the new row, do the following:
  - Click the **Router** field and then enter the IP address for the IPv4 or IPv6 seed router. Note that you can assign a seed IP address to different network views if your deployment has overlapping IP addresses.
  - Click the **Network View** field and then choose the network view you want to assign to the interface. A newly added seed IP does not have any associated network view by default.
  - Click the **Comment** field and then enter information about the seed router.

You can delete a seed router by selecting it and then click the Delete icon. Note that you cannot delete any seed router that is a default gateway.

### IPv6 Seed Router Usage

For the discovery of any IPv6 network, you must use seed router values, comprised of at least one well-connected IPv6 router, preferably with routes to all other networks to be managed. In some cases, seed routers may not have the full routing tables or be unable to provide full information for some reason. The general rule of thumb is that more seed routers are better, but the connectivity of seed router(s) also helps determine how many seed routers you need. Avoid having more seed router entries than necessary.



#### Note

For effective use of seed routers, provide SNMP credentials to the Probe member to allow it to pull the key routing and connectivity information, including the IPv6 routing table and the local Neighbor Discovery Cache, from the device. For more information, see [Defining Seed Routers for Probe Members](#) below.

### Configuring Advisor Properties

For information about Advisor, see [Monitoring Device Lifecycle and Vulnerabilities Using Advisor](#).

See the following pre-requisites:

- You have purchased the Advisor subscription.
- You have access to the internet, either through one of the Consolidator interfaces or through a proxy server.
- You have a Consolidator with the discovery service working on it.
- You have a local DNS resolver working on the discovery node. For more information, see [Enabling DNS Resolution](#).

To configure Advisor properties, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Discovery**.
2. Click **Edit > Grid Discovery Properties** in the Toolbar.

- Click the **Advisor** tab.

The screenshot shows the 'Infoblox (Grid Discovery Properties)' window with the 'Advisor' tab selected. The window is divided into several sections:

- Basic:**
  - Enable Advisor Application
  - \*Network Interface: Choose One (dropdown)
  - \*Execution Interval: 1 (input) Days (dropdown)
  - Execution Hour: 03 AM (dropdown)
  - Use proxy server
  - DNS Name or IP Address: (input)
  - \*Port: (input)
  - Credentials to connect to Proxy Server
  - \* Username: (input)
  - \* Password: (input)
- ADVISOR CENTRAL:**
  - \*API Endpoint Address: (input)
  - \*API Endpoint Port: (input)
  - Authentication:**
    - Token
      - \* Auth Token Value: (input)
    - Credentials
      - \* Username: (input)
      - \* Password: (input)
  - \*Minimum Severity: 8.0 (input)
  - Last Scheduled Execution Result: (input)
  - Last Run Now Result: (input)
  - Buttons: Test connection to central, Run Now

At the bottom of the window are 'Cancel' and 'Save & Close' buttons.

- Select **Enable Advisor Application**.  
Note this checkbox is available if a Consolidator exists in the Grid and the discovery service is working.
- Network Interface:** Specify one of the network interfaces of the Consolidator that runs Advisor. This interface is used for the internet connection to obtain the lifecycle and vulnerability data.
- Execution Interval:** Specify how often the Advisor service should be executed every N days or weeks.
- Execution Hour:** Specify the server hour when the Advisor service should be executed.  
Note the Advisor runs not at the exact hour specified, but at any minute of the specified hour.

8. If you do not want to expose your Grid member directly to the internet, select **Use proxy server**. Ensure that the proxy server has access to the internet.
9. Specify the following information for the proxy server:
  - DNS Name or IP Address
  - Port
  - Credentials to connect to Proxy Server (username and password)
10. Under **Advisor Central**, specify the following data:
  - **API Endpoint Address**: The IP address of the Advisor API endpoint.
  - **API Endpoint Port**: The port of the Advisor API endpoint.
  - **Authentication**: Select **Token** or **Credentials**.
  - If you selected token authentication, specify the authentication token value.
  - If you selected credentials authentication, specify the username and password.
11. In **Minimum Severity**, specify the severity threshold for vulnerabilities data that you want to obtain for your devices. To see possible values, hover the mouse over the field. The popup window displays the following values:
  - **Critical**: 9.0-10.0
  - **High**: 7.0-8.9
  - **Medium**: 4.0-6.9
  - **Low**: 0.1-3.9
  - **None**: 0.0
12. **Last Scheduled Execution Result**: Displays the timestamp of the last successful or unsuccessful scheduled execution result.
13. **Last Run Now Result**: Displays the timestamp of the last successful or unsuccessful immediate execution result.
14. **Test connection to central**: Central refers to the server where the application for Network Insight Advisor is running, that is the Advisor server. NIOS sends a POST query to the Advisor server from Discovery Consolidator, when you click **Test connection to central**. In the API Endpoint Address, the server address is specified. If the query enters the Advisor server successfully and returns a correct response, then OK is displayed else Not OK is displayed.
15. If you want to launch Advisor immediately, click **Run Now**.
16. Save the configuration.

## Configuring Discovery for SDN and SD-WAN

Network Insight allows you to collect and manage data from SDN and SD-WAN environments. Currently, you can discover Cisco ACI, Cisco Meraki, Cisco Viptela, and Juniper Mist. For information about adding, configuring, discovering data, and configuring polling settings for these devices, see:

- [Cisco Network Solutions](#)
  - [Adding and Configuring Cisco ACI Discovery](#)
  - [Adding and Configuring Cisco Meraki Discovery](#)
  - [Adding and Configuring Cisco Viptela Discovery](#)
- [Juniper Mist Network Solutions](#)
- [Configuring SDN and SD-WAN Polling Properties](#)

You can add specific SDN and SD-WAN entries in the discovery properties of Probe or Standalone members. You cannot configure these settings on Consolidators. Also, you can configure general SDN and SD-WAN polling properties in the Grid discovery settings.

To view discovery results for SDN and SD-WAN, go to **Data Management > Devices**. For information, see [Viewing Discovered Devices and their Properties](#).



### Note

To ensure successful SDN and SD-WAN discovery, use an admin user account.

## Cisco Network Solutions

You can add, configure, and discover data for the following Cisco solutions: Cisco ACI, Cisco Meraki, Cisco Viptela.

### Adding and Configuring Cisco ACI Discovery

Enabling the discovery of Cisco ACI devices provides visibility into your Cisco ACI infrastructure. This allows you to view and manage discovered IP addresses of Cisco ACI fabric members such as APIC controllers and fabric switches with their attached end points.



#### Note

The **Cisco APIC Configuration** tab in the member discovery properties was renamed to **SDN/SD-WAN**. You can find all previously configured Cisco ACIs in this tab that is described below.

To add and configure a Cisco ACI fabric, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Discovery**.
2. Select a Probe member, and then click **Edit > Member Discovery Properties** in the Toolbar.
3. Click the **SDN/SD-WAN** tab.
4. Click the Add icon and select **Cisco ACI**.
5. Complete the following:
  - **Fabric Name:** Specify a short and unique name for the current Cisco ACI configuration.
  - **Addresses:** Enter the hostname or IP address of the Cisco APIC controller. If your fabric includes more than one controller, click the Add icon to add more addresses.
  - **Protocol:** Select **HTTP** or **HTTPS**.
  - **Network View:** Select the network view to identify the corresponding network interface for connectivity with the Cisco ACI. Also, this network view will be assigned to discovered devices from this ACI.
  - **Username:** The login name for the Cisco ACI.
  - **Password:** The login password.
  - **Comment:** Additional information about the Cisco ACI.
  - **Connect using Grid Proxy settings if available:** Select if you want to use the Grid Proxy for connectivity to or from the Cisco ACI. If the Proxy is specified in the Grid properties, then Network Insight uses it. For more information, see [Configuring Proxy Servers](#).
6. Click **Test Connection** to check if the fabric is reachable and the provided credentials are correct. The connection test results are also written to syslog.
7. Click **Add**.
8. Click **Save & Close**.

### Adding and Configuring Cisco Meraki Discovery

Enabling the discovery of Cisco Meraki provides visibility into your Cisco Meraki SD-WAN elements, for example:

- Wireless access points
- Switches
- Routers
- Cameras
- Phones

Network Insight classifies Meraki cameras and phones as end hosts and other Meraki devices as network devices.

To add and configure Cisco Meraki discovery, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Discovery**.
2. Select a Probe member, and then click **Edit > Member Discovery Properties** in the **Toolbar**.
3. Click the **SDN/SD-WAN** tab.
4. Click the Add icon and select **Cisco Meraki**.
5. Complete the following:
  - **Config Name:** Specify a short and unique name for the current Cisco Meraki configuration.

- **Address:** Enter the hostname or IP address of the Cisco Meraki Dashboard API. By default, it is [api.meraki.com](https://api.meraki.com).
  - **Protocol:** HTTPS by default.
  - **Network Interface:** Select the interface that will be used to access the device.
  - **API Key:** An access key is required to use Cisco APIs.
  - **Comment:** Additional information about the Cisco Meraki device.
  - **Connect using Grid Proxy settings if available:** Select if you want to use the Grid Proxy for connectivity to or from the Cisco Meraki device. If the Proxy is specified in the Grid properties, then Network Insight uses it. For more information, see [Configuring Proxy Servers](#).
6. Click **Test Connection** to check if the device is reachable and the provided credentials are correct. The connection test results are also written to the syslog.
  7. Click **Add**.
  8. Click **Save & Close**.

### Adding and Configuring Cisco Viptela Discovery

Enabling the discovery of Cisco Viptela devices provides visibility into your Viptela SDN/SD-WAN infrastructure. You can use Viptela as an on-premises SDN controller or as a cloud solution.

To add and configure Cisco Viptela discovery, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Discovery**.
2. Select a Probe member, and then click **Edit > Member Discovery Properties** in the **Toolbar**.
3. Click the **SDN/SD-WAN** tab.
4. Click the Add icon and select **Viptela**.
5. Complete the following:
  - **Config Name:** Specify a short and unique name for the current Viptela configuration.
  - **Address:** Enter the hostname or IP address of the Viptela vManage controller.
  - **On-premise controller:** Check this if your Viptela setup is on-premises.
  - **Protocol:** Select **HTTP** or **HTTPS**.
  - **Network Interface:** Select the interface that will be used to access the device.
  - **Network View:** Select the network view in which the discovered Viptela devices will be shown.
  - **Username:** The login name for the Viptela vManage controller.
  - **Password:** The login password.
  - **Comment:** Additional information about the Viptela vManage controller.
  - **Connect using Grid Proxy settings if available:** Select if you want to use the Grid Proxy for connectivity to or from the Viptela. If the Proxy is specified in the Grid properties, Network Insight uses it.
6. Click **Test Connection** to check if the device is reachable and the provided credentials are correct. The connection test results are also written to syslog.
7. Click **Add**.
8. Click **Save & Close**.

### Juniper Mist Network Solutions

Enabling the discovery of Juniper Mist provides visibility into your Juniper Mist SD-WAN elements, for example:

- Wireless access points
- Switches
- Routers
- Firewalls

To add and configure Juniper Mist discovery, do the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Discovery**. The Grid Manager tab will be automatically chosen as the default selection.
2. Under the **Services** tab, select a Probe member. The **Edit** icon becomes active when you select the Probe member.
3. Click the **Edit** icon.
4. Click the **SDN/SD-WAN** tab from the left navigation pane.
5. Click the drop-down next to the Add icon.

6. Select **Juniper Mist**. The **Add Juniper Mist Configuration** page will be displayed.
7. Complete the following:
  - **Config Name:** Specify a short and unique name for the current Juniper Mist configuration.
  - **Address:** Enter the hostname or IP address of the Juniper Mist Dashboard API. By default, it is [api.mist.com](https://api.mist.com).
  - **Protocol:** This field is populated as **HTTPS** automatically and is not editable.
  - **Network Interface:** Select the interface for accessing the device.
  - **API Key:** An access key required to use Juniper Mist APIs.
  - **Comment:** Additional information about the Juniper Mist device.
  - **Connect using Grid Proxy settings if available:** Select if you want to use the Grid Proxy for connectivity to or from the Juniper Mist device. If the Proxy is specified in the Grid properties, then Network Insight uses it. For more information, see [Configuring Proxy Servers](#).
8. Click **Test Connection** to check if the device is reachable and the provided credentials are correct. The connection test results are also logged to syslog.
9. Click **Add**.
10. Click **Save & Close**.

## Configuring SDN and SD-WAN Polling Properties

The following devices are supported for SDN and SD-WAN polling settings:

- Cisco ACI
- Cisco Meraki
- Cisco Viptela
- Juniper Mist

On the Grid side, you can enable or disable the SDN and SD-WAN polling, specify end host collection timing, and define network view mapping rules. If SDN and SD-WAN polling is disabled, only traditional network devices are polled. Controlling the polling setting and end host data collection allows you to reduce the load on your system if required.

For the supported devices, you can select between different modes for mapping networks to NIOS network views. This mapping mechanism is required as your defined device infrastructure may have overlapping IP ranges that can be supported under different network views. The mapping rules include:

- Mapping to the predefined SDN network view
- Automatic mapping
- Custom mapping

To configure SDN/SD-WAN polling properties, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Discovery**.
2. From the Toolbar, select **Edit > Grid Discovery Properties**.
3. Click the **SDN/SD-WAN** tab.
4. On the **Basic** tab, complete the following:
  - **Enable SDN/SD-WAN polling:** Select to enable or disable SDN and SD-WAN polling.
  - **Default SDN Network View:** The network view that will be assigned to discovered Cisco Meraki devices for which the automatic network view mapping is disabled. You enable or disable automatic mapping in the advanced SDN and SD-WAN polling settings. For more information, see step 5 below.
  - **Detailed End Host Collection Interval:** Select to enable or disable the collection of end hosts (or clients in Cisco Meraki terminology). If enabled, specify one of the following:
    - **Periodic Collection:** Specify the N minutes or hours when the collection should occur.
    - **Scheduled Collection:** Schedule recurrent collection based on hourly, daily, weekly, or monthly time periods. Choosing this option, click the Calendar icon, and a Polling Scheduler appears; click the Edit icon to make scheduling changes. Choose a recurrence pattern of **Once**, **Hourly**, **Daily**, **Weekly**, or **Monthly**; in all cases, you must choose an **Execution Time**.
5. On the **Advanced** tab, complete the following:
  - **Disable SDN/SD-WAN Discovery for networks not in IPAM:** If set, new unmanaged networks discovered on the SDN controller are not created in the Infoblox IPAM.
  - **Network View Mapping:** Select one of the following:
    - **Disable automatic mapping and use predefined SDN Network View:** Select to map the collected SDN/SD-WAN devices to the default SDN network view defined in step 4 above.



- **Automatically create network views for unmapped networks:** Select to automatically map collected networks to their network views using Network Insight's internal rules. Network views that do not exist are created automatically. The mapped networks are displayed in the table that is not editable.
- **Enable network view mapping defined below:** This is custom mapping. Select to manually map the collected networks to appropriate network views. To change a network view entry, click it in the table.

6. Click **Save & Close**.



#### Note

A network name in the mapping table is made up by combining a vendor name and network name. The **Source** column displays the fabric name or config name that you previously defined for the SDN or SD-WAN configuration. The name of a network view is made up by combining a network value with a source value.

## Consolidator and Probes

When you first join a discovery member to the Grid, the first discovery member that joins the Grid automatically becomes the Consolidator and all other discovery members become Probes. If you have only one discovery member in the Grid, it becomes the Consolidator-Probe standalone discovery appliance.

### Consolidator

The central repository of discovery data for the entire managed network. The Consolidator is a single appliance that contains data about all devices detected through discovery. The Consolidator communicates with the Grid Master as a normal Grid Member and transfers all its data to the Grid Master, as indicated in the figure [Network Insight Appliances Added as Grid members](#). The Consolidator compiles information from one or more associated Probe appliances. The Consolidator appliance requires the Discovery license. If you have one or more Probe appliances or virtual appliances, the Consolidator performs no discovery on its own. If you plan to use a single dedicated appliance for discovery, that appliance must be licensed for discovery and be configured as a Consolidator. Note that the Grid Master cannot be licensed as a Consolidator.

### Probes

A Probe is a Network Insight appliance or virtual appliance that performs the direct querying, probing and polling of network devices and the initial data collection. Probe appliances also require the Discovery license. Infoblox recommends using one or more Probe appliances with the Consolidator. Each Probe can override the Grid level discovery credentials with its own discovery credentials. Data synchronization occurs continuously between the Consolidator and all associated Probe appliances and between the Consolidator and the Grid Master.



#### Note

You assign each Probe appliance to a single network view, and multiple Probe appliances can share the same network view. You can change network view assignments for Probe appliances at any time. On ND-1405, ND-2205, ND-4000, ND-V1405, and ND-V2205 Network Insight appliances, you can assign multiple VLAN interfaces on the same Probe to different network views.

## Consolidator Probe Appliance

You may also choose to operate a Consolidator-Probe appliance as a single discovery system. In this deployment, the appliance operates as both a Consolidator and a Probe, performs all discovery operations, aggregates all databases within it, and synchronizes with the Grid Master.

Standalone discovery appliances cannot be installed in a network that already has existing Probes and a Consolidator. For more information, see [Defining the Discovery Member Type](#) and [Mapping Discovery Interfaces to Network Views](#).

## Deployment Guidelines for Consolidator and Probes

When you wish to install and deploy discovery appliances, use the following installation guidelines:

- **Installing a Standalone in the Grid**—Before you designate an appliance as a standalone discovery appliance, no previously installed Probes should be present on the network and joined to the Grid. If you install a new appliance intended as a standalone, in a network that already has one or more Probe instances (perhaps for testing or evaluation purposes), before discovery service is stopped on the Probe instances, the new "standalone" appliance automatically detects the Probe instances and start as a Consolidator appliance, preventing it from acting to probe and detect devices as a standalone appliance. Consolidators cannot be assigned to network views or to discovery in network objects such as IPv4 or IPv6 network containers.
- **Converting a Consolidator to a Standalone**—Also consider the example of a Consolidator appliance operating with one or more instances running as Probes, each with respective Discovery licenses. If you wish to convert the Consolidator to a standalone discovery appliance, stop the discovery service on all associated probes. Then, stop and restart the discovery service on the Consolidator appliance. The appliance is selectable for discovery of network objects, acting as a standalone discovery appliance.
- **Adding new Probe Instances to a Standalone deployment**—Finally, consider the use of a standalone discovery appliance to which you wish to associate a new Probe instance or instances. This process converts a standalone to a Consolidator. After the new Probe instances join the Grid, stop the discovery service on the standalone discovery appliance. Then, start the discovery service on the new Probe or Probes. Next, restart the discovery service on the previously defined standalone appliance. It detects the newly active Probe instances and activate as a Consolidator.
- In all cases, you must maintain proper Discovery licensing.

## Defining the Discovery Member Type

Before using Network Insight to discover devices and networks, you must first define the consolidator and probes in your Grid and specify the discovering interfaces on these members.

You choose between the following member types:

- **Probe:** Turns the appliance into a discovery Probe appliance.
- **Consolidator:** Turns the appliance into a discovery Consolidator appliance.
- **Unassigned:** Disables the discovery feature on the appliance.

When you first join discovery members to the Grid, the first discovery member that joins the Grid automatically becomes the Consolidator and all other discovery members become Probes. This is the Grid mode. If you have only one discovery member in the Grid, it becomes a Consolidator-Probe discovery appliance – this is the standalone mode.

On some occasions, you may wish to change an appliance with a Discovery license to a Consolidator or to a Probe, or change a Consolidator to a standalone discovery appliance. To make any of these changes, you must first stop the discovery service on the appliance. For more information, see [Starting and Stopping the Discovery Service](#).

To define the discovery member type:

1. From the **Grid** tab, select **Grid Manager** -> **Discovery** to display the list of members running the discovery service.
2. Select the discovery member for which you wish to change the member type.
3. If the discovery service is running on the member, stop it.
4. In the Toolbar, click **Edit** -> **Member Discovery Properties**.
5. In the **General** tab -> **Member Type** field, choose the required member type.

Note that the **Consolidator** and **Unassigned** member types in the Grid mode are disabled if the member has SDN/SD-WAN configurations on it. This is because SDN/SD-WAN configurations can be added to only Probe or standalone member types. In the standalone mode all member types are available for selection.

6. Save the configuration.

To change the interface over which a member sends and receives discovery traffic, see [Mapping Discovery Interfaces to Network Views](#).

## Viewing Discovered Devices and their Properties

Grid Manager provides a **Devices** page under **Data Management** for a complete list of every device that discovery finds, and lists all unmanaged and managed devices. Here you can explore information about the discovered devices and drill down to specific information about every device. For information, see the next section.

Listed devices can be displayed in one of three states in the **Devices** page:

- Devices that appear with an empty value in the **Managed** column are devices that are discovered, but are not recognized by IPAM, are not part of an IPAM network, and hence cannot be changed to managed status in Grid Manager. These discovered devices cannot be changed to managed status, but you can perform actions such as activating and deactivating ports, executing **DiscoverNow** on the device, view their list of connected networks, and other actions. Avoid changing the state of ports or taking other actions on a discovered device, unless the action is verified by an administrator.
- Device shown in yellow table rows are unmanaged devices, but are recognized by IPAM and can be converted to managed status. Yellow rows appear with a value of **No** in the **Managed** column. You can convert devices in yellow table rows to managed objects under IPAM (host, fixed address, A record or PTR record).
- Devices shown in light grey table rows are managed devices, with a value of **Yes** in the **Managed** column.



### Note

For information about managed and unmanaged devices, see [Converting Unmanaged Devices to Managed Devices](#). You can also use the "Unmanaged devices and networks" filter in global search to locate all the unmanaged devices and networks discovered through discovery. For more information, see [Using Global Search](#).

Also, you can view VRF-based devices and map them to network views from the **Data Management** → **Devices** tab. See [Viewing Discovered VRFs and Mapping Network Views](#) below.

## Viewing the Complete List of Discovered Devices

The **Data Management** tab → **Devices** tab provides a complete view of all discovered devices discovered by Network Insight. The list includes routers, switches, firewalls and other security devices, wireless APs, end hosts and servers in end-host networks. Use NIOS standard filtering to narrow down the status table to the devices or values you want to examine.

You can see the following information in the devices table:

- **IP Address:** The detected management IP address (IPv4 or IPv6).
- **Name:** Detected name of the device. Each device name provides a link to the complete body of information associated with the device, arranged in five tabs: **Interfaces**, **Networks**, **IP Addresses**, **Assets** and **Components**. For more information, see the sections under [Accessing Detailed Device Information](#).
- **Device Type:** The network device type: **Router**, **Switch-Router**, **Firewall**, **NIOS (Infoblox appliance)**, **vNIOS**, **SDN Controller**, **SDN Element**, **LWAP**, and others.
- **Model:** The model name as detected by the device during discovery.
- **Serial Number:** The serial number of the discovered device.
- **Vendor:** The equipment manufacturer (Cisco, Juniper, Fortinet, F5, and others).
- **Device Version:** The Operating System version for the network device.
- **Chassis S/N:** The chassis serial number of the discovered device.

- **Location:** The physical location of the network device as detected by the device during discovery.
- **Description:** Verbose description of the network device as collected from the device by discovery.
- **Discover Now:** Indicates when the device is undergoing a current discovery process. A "Pending" icon appears in this column to indicate the status.
- **Managed:** Indicates the status of the device in Grid Manager. A blank value in this field indicates the device has been discovered but is not recognized in IPAM; a **No** value indicates the device is recognized by IPAM but is not managed under Grid Manager; and a **Yes** value indicates that the device is fully managed by Grid Manager from use of the **Convert** command and can support related features such as port reservations and IPAM/DHCP object assignments.
- **Active Users:** The number of active users on the Active Directory domain for the selected IP address.

For each listed device, the Action icon  provides the following options depending on the device type and its status:

- **Show IPAM IP Address:** Shows the management IP address for the device that has a network in IPAM—the main **IPAM** tab appears, showing details for the IP address. This option is greyed out for devices that have a management IP that is not part of an IPAM network.
- **Edit:** Displays the Device Properties Editor window. Alternatively, you can select the device and click the Edit icon above the devices table. For more information, see [Editing Device Properties](#) below.
- **Interfaces:** A direct link to the **Interfaces** page for the chosen device. Unmanaged devices may have managed interfaces that appear in this page, and managed devices may have unmanaged interfaces that appear here. For more information, see [Viewing Interface Information for Discovered Devices](#).
- **Discover Now:** Immediately performs discovery on the selected device.
- **Convert:** For devices in unmanaged status (shown in yellow), allows conversion of the device to a managed object in Grid Manager: a host, fixed address, A record or PTR record. For more information, see [Converting Unmanaged Devices to Managed Devices](#).
- **IPAM Networks:** A drop-down list of all IPv4/IPv6 IPAM networks currently provisioned on the device. Each network provides a link to the IP Map page for the selected network.
- **Device Details:** A basic list of information about the chosen device, including the IP address by which the device is discovered, operational status, IPAM Type (whether the device is managed or unmanaged), the Device Type and the number of Interfaces.
- Click **Device Support** to verify data collection activities in the following tabs:
  - **Data Collection:** You can view the timestamp at which the most recent collection from various data sources was completed. The sources from which device support information is collected are listed under the Data Source column, and it includes the device's routing table (ipRouteTable), environment monitoring (DeviceEnvMon), and numerous other data sources as applicable to the specific device type. It displays the following information for each discovered device:
    - **Data Source:** The sources from which the device support information was collected.
    - **End Time:** The most recent timestamp of the data collected by the discovery member.
  - **Device Support:** Lists various types of information supported for collection on the current device. You can view the following details for each discovered device:
    - **Function:** Data function that can be collected by Network Insight. The value can be **Device Vendor**, **Device Model**, **Device Version**, **VLANs**, **Forwarding**, **VRFs**, **Inventory**, and **SecurityControl**.
    - **Supported:** Indicates whether this data function is supported for the selected device. The value can be **Yes** or **No**. If it is **No**, Network Insight will not attempt to gather the data. For instance, for a Cisco router, Network Insight does not attempt to gather VLAN information, so a **No** value will be displayed in the **Supported** column.
    - **Available:** Reflects whether the data has actually been collected. The value can be **Yes** or **No**. A value of **Yes** for **Supported** and **No** for **Available** indicates a discovery misconfiguration or could possibly require an adjustment to the Device Support Bundle (DSB) for that particular device model.
    - **Value:** Displays the value collected for the **Device Vendor**, **Device Model**, and **Device Version** data functions. Displays **Last Collected** time for the **VLANs**, **Forwarding**, **VRFs**, **Inventory**, and **Security Control** data functions.
- **Show Active Users:** Displays the *Active Users* dialog box. You can view all the active users on the Active Directory domain for the selected device. For more information, see [Viewing Active Network Users](#).

Click **Discovery Status** in the Toolbar to view the same list of network devices showing the discovery data set. You can sort the table by Name or IP address. Use Grid Manager-standard filtering to display device names, IP addresses or other values in which you are interested.

## Editing Device Properties

In the Device Properties Editor, you can change the management IP address and settings for interfaces of the selected device, apply extensible attributes, or apply administrative permissions for Grid Manager admin access to the device.

To edit properties of a device:

1. In the **Data Management** → **Devices** tab, click the Action icon  
for the required device or select the device and click the Edit icon above the table.
2. In the **General** tab, edit the following general device properties:
  - **Name:** The discovered device name, such as SPINE, LEAF, switch1.building2.com, or office1router.
  - **Management IP address:** This IP address is used for the device to be discovered by the discovery member. If the device has more than one IP address that can be used as the management IP, you can manually select the required address from the list of those discovered on the device.  
If more than one network view is used in the discovery, the name of the corresponding network view is displayed next to the IP address.  
Note changing of management IP address may take some time.
  - **Managed:** Indicates whether this device is managed or unmanaged in NIOS.
  - **Device Type:** Indicates the device type: Router, Switch-Router, Firewall, NIOS (Infoblox appliance), vNIOS, SDN Controller, SDN Element, and others.
  - **Description:** Description of the network device as collected from the device by discovery.
  - **Model:** The device model.
  - **Vendor:** The vendor of network device. For example, Cisco, Juniper, Aruba, Dell, Infoblox, or HP.
  - **Device Version:** The OS version of the device.
  - **Location:** Displays the physical location of the device as detected during discovery.
  - **Discover Now:** Displays the status of the Discovery Now operation of the device.
  - **Object Type:** The object type.
3. In the **Interfaces** tab, edit the following basic port settings:
  - **Admin Status**
  - **Description**
  - **Data VLAN**
  - **Voice VLAN**
4. In the **Extensible Attributes** tab, add any attributes that are necessary for the device.
5. In the **Permissions** tab, edit administrative permissions for the device:
  - a. Click the Add icon. Grid Manager displays the Admin Group/Role Selector dialog box.
  - b. In the Admin Group/Role Selector dialog box, select the admin group you want to add, and then click the Select icon.
  - c. Select the permission for the group.
  - d. Select an object to which the permission applies from the drop-down list in the Resources column.  
Note that administrators can access discovered and managed devices in Grid Manager. For tasks such as provisioning networks, adding administrative permissions is advised to ensure that unauthorized changes to device configurations cannot take place. For example, you can use accounts with the Port Control permission to control and manage network provisioning tasks.
6. Click **Verify** where applicable, for example, in the interfaces settings.
7. Click **Save & Close**.

## Viewing Discovered VRFs and Mapping Network Views

To view VRF instances (or VRFs) and map corresponding network views, do the following:

1. From the **Data Management** tab, select the **Devices** tab, and then click **VRF Mapping** from the Toolbar.
2. The *VRF Mapping dialog* appears and displays the following:
  - **VRF Name:** The name of the VRF on the hosting device, which typically contains the interface name and its VRF route distinguisher.
  - **Device Name:** The discovered name of the device that is hosting the VRF.
  - **Device IP Address:** The IP address of the managed VRF hosting device.
  - **Network View:** The network view that is associated with the VRF. You can click this field and select a different network view from the drop-down list.

You can do the following in this tab:

- To assign the same network view to multiple VRFs, select the checkboxes of the VRFs, and then click the Edit icon. The *VRF Mapping* dialog displays the **Edit VRF Network View** panel. From the **Network View** drop-down list, select the network view you want to assign to all the selected VRFs, and then click **Save**. If there is only one network view in the Grid, which is the default view, the **Network View** column is hidden by default.
- You can use filters to narrow down the list. You can filter the list based on the VRF name, Device name, Device IP address, and network view. For more information, see [Using Filters](#).
- You can sort the data in ascending or descending order by column.



#### Note

The appliance displays a warning message when there are discovered VRFs that are not mapped to network views. To ensure that discovered VRFs are mapped to network views, you can configure automatic VRF mapping, as described in [Configuring Automatic VRF Mapping](#).

## Port Control Features in Network Insight



#### Note

Port control involves two primary operations: *network provisioning/de-provisioning* and *port configurations*. These operations are classified as port control tasks that can be monitored and viewed in the Task Manager (**Administration** → **Workflow** → **Task Manager**).

Port control enables changes to the interface-level configurations of switches and switch-router devices, and assignment of these resources to network objects defined and created within IPAM.

- Port configurations and network provisioning and de-provisioning use CLI admin credentials, supporting SSH and Telnet. You may test credentials before use, against an IP address or a selected device;
- Port configuration consists of two primary operations: setting admin status for a port, and defining Data VLAN and Voice VLAN assignments (where applicable), along with minor changes such as editing descriptions;
- You can define port configuration blackout periods using the same methods provided for discovery blackouts. These blackout periods also apply to network provisioning and de-provisioning tasks;
- Configuring a port on a device always creates a new port control task that can be viewed and managed in the Task Manager.
- You can separately schedule port control tasks using the same method as for object creation.
- You can edit multiple interfaces at a time.
- You can edit interfaces, inline, from the **Interfaces**, **IP Addresses** and **Assets** pages in Grid Manager. These operations generally consist of setting the interface to be Administratively Up or administratively Down, and VLAN assignments.

Network provisioning includes the following:

- If a user deletes a discovered network from the system, Grid Manager displays the list of *interfaces* on which the network is currently provisioned;
- Network provisioning allows you to provision a network on one interface at a time. The network must be in managed status under IPAM;
- The user can also *de-provision* a network, which removes it from one or more interfaces;
- You can perform network provisioning and de-provisioning tasks on routers and switch-routers.

Devices do not have to be in managed state for some port control operations (setting ports to Admin Up and Admin Down, for example) but some port control operations require it:

- Provisioning a network (through IPAM) onto a port on a managed device;
- De-provisioning a network.

When you create a new object using the wizard, you can configure the port or ports that are associated with the object's port reservation. In this case, two new tasks are created: an object creation task, and a port control task, which can be scheduled separately from the object creation. *The port control task is a separate task that may also require administrator approval.* When you create a new task, an information feedback panel provides a link to the port control task in the Task Manager.

You may also select and reschedule both tasks. For more information, see [Rescheduling Tasks](#).



#### Note

If you edit an object, you can only edit an associated port reservation.

Objects are completed in their configuration by Grid Manager before executing a port configuration.

If, for example, a fixed address object is subject to administrative approval, no port control task takes place for that object until the approval is executed and the object is created. This has implications for scheduling: if you schedule the creation of a new host, IPv4 reservation or fixed address, and wish to schedule a port control task for the same object, the scheduled object creation must take place first, and must complete, before the scheduled port configuration executes.

All port configuration operations can be scheduled and subject to administrator approval. For more information, see [Configuring Approval Workflows](#).

## About Port Reservations

You can define a device *port reservation* for a defined object such as a host, fixed address or IPv4 reservation, or for an Infoblox Grid Member. Port reservations assign device interfaces to Infoblox-managed objects in Grid Manager. The port reservation is a property of each object, such as a fixed address or host record, that you create. After discovering and cataloguing infrastructure devices, Grid Manager compiles the lists of interfaces and separately tracks available switch ports on each device. When you attempt to reserve a switch port, Grid Manager provides the complete list of *available* switch ports on the device, automatically preventing possible conflicts over port usage.

Characteristics of port reservations include the following:

- Device Ports can be reserved to Grid Manager objects under IPAM and DHCP;
  - The user can choose to immediately create the new object (IPv4 reservation, IPv4/IPv6 fixed address or host, or a Grid member), and to also immediately assign a device port to the object. In such cases, the object is created and the port is reserved for the object;
  - An object can be scheduled for creation at a later time, and its port assignment scheduled for the same time, or for a different date and time, in the future;
  - A port reservation does not guarantee that the switch interface is in fact available for the assignment.
- You can create port reservations for multiple objects at a time;
- Interface tables show the objects to which they are bound, in a special **Reservation** column;
- When you delete an IPAM object such as a fixed address, its associated port reservation is automatically deleted.

Once a switch port or other device port is reserved, Network Insight prevents future tasks from trying to use the same port for another reservation.

Device switchports and router interfaces, and other interface types may be assigned to the following types of objects:

- Grid members (including HA Pairs). For more information, see the following sections [Defining Port Reservations for an Infoblox Grid Member](#) and [Defining Port Reservations for an HA Pair](#).
- Hosts. For more information about defining hosts with included port configuration, see [Adding Host Records](#).
- IPv4 reservations. For more information, see [Adding IPv4 Reservations](#).
- Fixed addresses (IPv4 and IPv6). For more information, see [Adding IPv4 Fixed Addresses](#) and [Adding IPv6 Fixed Addresses](#).
- IP networks (IPv4 and IPv6). For more information, see [Adding IPv4 Networks](#) and [Adding IPv6 Networks](#)

Devices involved in these operations must be under managed status in Grid Manager. For more information, see [Converting Unmanaged Devices to Managed Devices](#).



## Defining Blackout Periods



### Note

You can separately define blackout periods for discovery, and for port configuration. This section describes how to use the blackout feature for discovery. For more information on blackouts for port configuration tasks, consult the section [Defining Port Configuration Blackout Periods](#).

Discovery protocols can occupy significant resources within the network when discovery is taking place. While you can schedule any discovery or port control task for any single time period or recurring time period, you can also establish time periods when Network Insight does not talk to devices or networks for discovery.

You can define blackout settings at two levels in Grid Manager:

- Under Grid Discovery Properties, applying across the entire Grid;
  - All networks managed by the Grid inherit discovery blackout settings by default;
- For individual networks under IPAM and under DHCP.
  - A network must be Managed before you can edit its discovery blackout settings.
  - Under IPAM, you can define discovery blackout settings for Network Containers and for networks (for DHCP, you can also set blackouts for DHCP Ranges);
  - If a network is in Managed state, it can be edited under IPAM or under DHCP for discovery settings and discovery blackout settings.

Discovery tasks may already be running when a blackout period takes effect. Current tasks are not interrupted and will complete within their time. Network Insight does not activate new discovery tasks during the blackout period, however.

### Defining Blackout Periods for the Grid

This procedure also covers defining port configuration blackout periods for the Grid. To define blackouts for the Grid:

1. From the **Grid** tab, select the **Grid Manager** tab, and click **Edit → Grid Discovery Properties**.  
If you do not select an appliance from the list, the **Grid Discovery Properties** option remains enabled.
2. Expand the Toolbar and click **Edit → Grid Discovery Properties**.
3. Click the **Blackout** tab.
4. Select the **Enable Discovery Blackout** checkbox and click the Scheduling icon to open a separate scheduling window. The Blackout Scheduler dialog opens.
  - a. Select how often you want to execute the blackout period. You can select **Once**, **Daily**, **Weekly**, or **Monthly**.
  - b. If you select **Once**, enter the day in the date picker and select a month from the drop-down list.
    - Enter a time in the hh:mm:ssAM/PM format. You can also select a time from the drop-down list.
    - Choose the **Time Zone**.
    - Specify the **Duration**: 1 or more Minutes, Hours or Days.
  - c. If you select **Daily**, click either **Every Day** or **Every Week day**.
    - Enter a time in the hh:mm:ssAM/PM format. You can also select a time from the drop-down list.
    - Choose the **Time Zone**.
    - Specify the **Duration**: 1 or more Minutes, Hours or Days.
  - d. If you select **Weekly**, complete the following:
    - Under **Schedule every week on:**, select the checkbox for any day of the week.
    - Enter a time in the hh:mm:ssAM/PM format. You can choose a time from the drop-down list.
    - Choose the **Time Zone**.
    - Specify the **Duration**: 1 or more Minutes, Hours or Days.
  - e. If you select **Monthly**, complete the following:
    - **Schedule the day of the month**: A discovery blackout can be executed monthly on a specific day, or instances can be executed more than one month apart on a specific day, in the **Day every month(s)** field.
    - Enter a time in the hh:mm:ssAM/PM format. You can choose a time from the drop-down list.
    - Choose the **Time Zone**.
    - Specify the **Duration**: 1 or more Minutes, Hours or Days.



5. If necessary, select the **Enable Port Configuration Blackout** checkbox and click the Scheduling icon to open the scheduling window. For more information, see [Defining Port Configuration Blackout Periods](#).
6. Follow Steps 5a–5e to schedule the port configuration blackout for the Infoblox Grid.  
–or–  
Select the **Use Discovery Blackout Schedule** checkbox to apply the discovery blackout schedule to port configurations.
7. When you have finished configuring schedules for blackout periods, click **Save & Close**.

When a scheduled Grid-level blackout period goes into effect, no operations related to discovery take place for the specified time period across the Infoblox Grid. Any discovery tasks already in progress will run to completion but no new ones will start.

## Defining Blackout Periods for Networks

Network Insight offers considerable flexibility in how you apply blackout periods. You may choose to have discovery allowed for most managed networks but elect to have a discovery blackout for selected networks that are traffic- or latency-sensitive.

You can define extended time periods and regularly scheduled times when discovery and/or port configuration tasks is not in progress on a specific IPAM network or within a network container. By default, the network inherits its discovery blackout settings from the Grid level. Editing a network under IPAM or DHCP, blackout settings apply only to the specified network. You also specify the scheduled time when the blackout period begins and the duration of the blackout period. As noted, a network must be in managed status before editing discovery or blackout features. To define a discovery blackout for a network under IPAM or DHCP:

1. Select a managed network from one of the following locations:
  - a. **Data Management → IPAM → list view**  
–or–
  - b. **Data Management → DHCP → Networks**
2. Click the Action icon  
  
next to the network you want (this automatically selects it) and select **Edit** from the menu. The Edit Network dialog appears.
3. Click the **Discovery Blackout** tab.
4. Click **Override** to change blackout settings for the chosen network.
5. Select the **Enable Discovery Blackout** checkbox and click the Scheduling icon to open a separate scheduling window. (Because the settings are inherited, **Enable Discovery Blackout** may or may not already be enabled.) The Blackout Scheduler dialog opens.
  - a. Select how often you want to execute the blackout period. You can select **Once**, **Daily**, **Weekly**, or **Monthly**.
  - b. If you select **Once**, enter the day in the date picker and select a month from the drop-down list.
    - i. Enter a time in the hh:mm:ss AM/PM format. You can also select a time from the drop-down list.
    - ii. Choose the **TimeZone**.
    - iii. Specify the **Duration**: 1 or more Minutes, Hours or Days.
  - c. If you select **Daily**, click either **Every Day** or **Every Weekday**.
    - i. Enter a time in the hh:mm:ss AM/PM format. You can also select a time from the drop-down list
    - ii. Choose the **TimeZone**.
    - iii. Specify the **Duration**: 1 or more Minutes, Hours or Days.
  - d. If you select **Weekly**, complete the following:
    - i. Under **Schedule every week on**: Select the checkbox for any day of the week.
    - ii. Enter a time in the hh:mm:ss AM/PM format. You can choose a time from the drop-down list.
    - iii. Choose the **TimeZone**.
    - iv. Specify the **Duration**: 1 or more Minutes, Hours or Days.
  - e. If you select **Monthly**, complete the following:
    - **Schedule the day of the month**: A discovery blackout can be executed monthly on a specific day, or instances can be executed more than one month apart on a specific day, in the **Day every month(s)** field.
    - Enter a time in the hh:mm:ss AM/PM format. You can choose a time from the drop-down list.
    - Choose the **Time Zone**.
    - Specify the **Duration**: 1 or more Minutes, Hours or Days.

6. If necessary, select the **Enable Port Configuration Blackout** checkbox and click the Scheduling icon to open the scheduling window. For information, see [Defining Port Configuration Blackout Periods](#).
7. Follow Steps 5a–5e to schedule the port configuration blackout for the chosen network.  
–or–  
Select the **Use Discovery Blackout Schedule** checkbox to apply the discovery blackout schedule defined for the network.
8. When you have finished configuring schedules for blackout periods for the network, click **Save & Close**.

When a scheduled blackout goes into effect, no operations related to discovery and/or port configuration take place for the specified time period on the selected network. Any related operations in progress will run to completion but no new ones will start.

## Defining Port Configuration Blackout Periods

Similar to discovery blackout periods, you can define port configuration blackout periods for managed networks under IPAM and DHCP. You specify the scheduled time when the blackout period begins, and the duration of the blackout period.

By default, networks inherit their blackout settings from the Grid level. For port configuration blackout settings in a network, the network must be under managed status in IPAM.

### Defining Port Configuration Blackouts for the Grid

Port configuration blackout settings apply globally across the entire Infoblox Grid unless overridden by members or for specific networks.

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and click **Edit → Grid Discovery Properties**.
3. Click the **Blackout** tab.
4. Click the **Enable Port Configuration Blackout** checkbox.
  - a. Select the **Use Discovery Blackout Schedule** checkbox to apply the same blackout schedule to port configurations that is applied locally to discovery blackouts. (For more information on discovery blackouts, see [Defining Blackout Periods](#).)  
–or–
  - b. Click the Schedule icon below. The Blackout Scheduler dialog opens.

### Scheduling a Port Configuration Blackout

The screenshot shows the 'Blackout Scheduler Dialog' window. It features a title bar with a close button. On the left side, there are four radio buttons for scheduling frequency: 'Once' (selected), 'Daily', 'Weekly', and 'Monthly'. The main content area is titled 'Schedule once' and contains the following fields:

- Start Date:** 2019-01-08
- Start Time:** 02:40:07 PM
- Time Zone:** (UTC - 8:00) Pacific Tir
- Duration:** (empty field) Minutes

At the bottom of the dialog, there are 'Cancel' and 'OK' buttons.

## Defining Port Configuration Blackouts for Networks

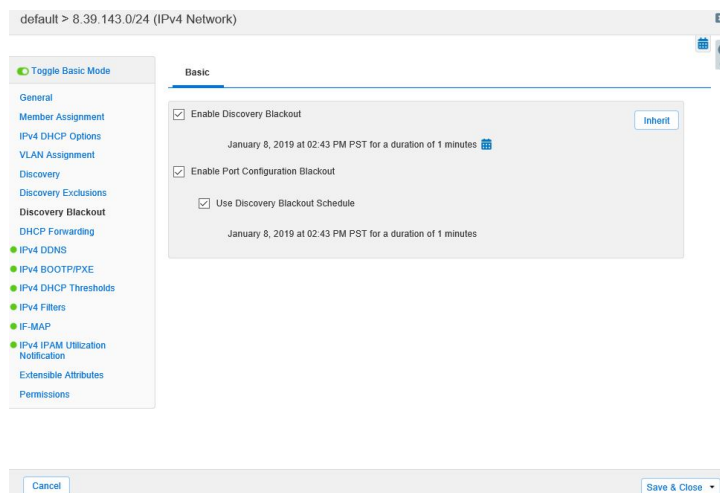
By default, networks inherit their blackout settings from the Grid level. You can override these with local settings for the network or for a network container, for both IPAM and DHCP. To change port configuration blackout settings for an IPAM or DHCP network or network container:

1. Select a managed network or network container from one of the following locations:
  - a. **Data Management** → **IPAM** → list view
  - or–
  - b. **Data Management** → **DHCP** → **Networks**
2. Click the Action icon

for a chosen network and choose **Edit** from the popup menu.

- Under IPAM, the network must be in managed status (a value of **Yes** appears in the **Managed** field, and the table row is highlighted in white).
  - Under DHCP, all networks appearing on this page may be selected for this purpose.
3. Click the **Discovery Blackout** tab. The editor page appears as shown in the below figure.

*The Discovery Blackout tab with enabled Port Configuration Blackouts*



## Creating Port Reservations for IPAM Objects

A port reservation instructs Network Insight to reserve ports on discovered and managed devices, for exclusive use by Grid members and by IPAM objects such as hosts or fixed addresses. Network Insight ensures that doing so does not interfere with existing device and port configurations and active networks, because port reservations automatically apply only to ports that are discovered to be available on network devices. Network Insight prevents a user from reserving the same port for more than one object. Should a port reservation somehow conflict with another more recent port reservation, Network Insight automatically reports a conflict and enables you to respond to the issue (for more information on this topic, see [Resolving Port Reservation Conflicts](#)).

When you create new IPAM objects, you can create the new objects and define port reservation settings at that time, or create the new object without any port-related settings and edit them later, after the object is established in Grid Manager. The Add IPv4 Reservation Wizard, Add Fixed Address Wizard, Add Host Wizard and Add Grid Member Wizards all support the full set of port configuration settings, as shown in the below figure.

*Defining Port Reservations during Object Creation*

Add IPv4 Fixed Address Wizard > Step 6 of 8

Reserve Port Reserving a Switch Port does not ensure its availability

\*Device Name: QAalabswrtr04.net.com Select Device Clear

\*Interface: Gi1/0/1

The following VLANs are configured:

ID	Name
1	default

Configure Port

Voice VLAN: Choose One

Data VLAN: 20 VLAN0020

\*Admin Status: Up

Description:

Cancel Previous Next Schedule for Later Save & Close

Making a port reservation does not guarantee that the port is in fact available on the requested device. All interfaces appearing in the Interface list are ports that are otherwise known to Network Insight as Operationally Down during its last discovery task, and that are not already reserved by a port reservation.

1. Begin by checking the Reserve Port checkbox.  
Note that optionally, you can completely skip port reservation and port configuration by clicking **Next**.
2. Click the Device Name button to choose the device for which the port reservation is associated. You should know the identity of the device to which the Infoblox appliance is connected before taking this step.
3. After choosing the device, choose the Interface with which the reservation is bound. The drop-down list shows only interfaces that are most recently found to be available by Network Insight during the last discovery cycle. This list does not include any ports that are Administratively Up and Operationally Up and are otherwise already assigned to other networks or objects.
  - The Wizard page also shows a list of any VLANs that are currently configured in the chosen device. This Wizard page does not allow the definition of new VLANs for port configuration—only the assignment of an existing VLAN in the device for port configuration.
4. Select the Configure Port checkbox to define specific port control settings for the port reservation.  
Note that if you do not take this action when you create the object, you cannot perform the configuration later while editing the object.
5. If the chosen device supports them, choose the Voice VLAN and/or the Data VLAN settings you may need for the port assignment. You do not create new VLAN values in this step; you can select from VLANs that are provisioned on the currently chosen device. All VLANs configured on the device and discovered by Network Insight during its most recent discovery polling cycle appear in the drop-down lists.
6. Set the Admin Status to Up if you need to activate the port in the current task. Though the port reservation is associated only with the current object, any port configuration creates a new Port Control task under Task Manager.
7. Enter a description for the port assignment. Infoblox recommends doing so to help other technicians to recognize the port assignment event.
8. When finished, click **Save & Close** or select other tabs to change settings for the object.  
Note once a switch port or other device port is reserved, Network Insight prevents further port reservations from using the same port for another reservation.  
See the following sections for examples on how to create IPAM objects with port reservations:

- [Adding Host Records](#)
- [Configuring IPv4 Networks](#)
- [Configuring IPv4 Fixed Addresses](#)
- [Configuring IPv4 Reservations](#)
- [Configuring IPv6 Networks](#)
- [Configuring an IPv6-only Grid](#)
- [Adding Grid Members](#)
- See Defining Port Reservations for an Infoblox Grid Member below

### Editing Port Reservation Settings for IPAM Objects

As previously noted, you can create IPAM objects such as hosts or fixed addresses without device information and port reservation settings, and edit them later after the object is established in the Grid Manager. You can change port reservations in any object to new settings. Limitations exist when editing existing objects.

The **Port Reservation** editing page settings uniformly apply to Grid members, IPv4 reservations, hosts, and fixed address objects. Note that you cannot change port configurations when editing objects, as shown in the below figure.

#### Editing an Existing Object's Port Reservation

192.168.1.1 (Fixed Address)

Toggle Advanced Mode

Basic

Reserve Port Reserving a Switch Port does not ensure its availability.

\*Device Name QA1abswrtr04.net.com Select Device Clear

\*Interface Gi1/0/1 To configure this interface, go the Devices tab and open the Interfaces tab for the selected device.

The following VLANs are configured:

ID	Name
1	default

Cancel Save & Close

Unlike creating an object, editing an existing object's port reservation does not permit configuring the selected port. (You can edit the port from the device's **Interfaces** page, including inline editing.) Physical interfaces with an Operational Status of Down appear in the **Interface** drop-down list. Ports that are already active, that are reserved through a port reservation, or that are administratively Up/Operationally Up do not appear in the **Interface** drop-down list.

For object editing, you can select interfaces but you cannot edit their settings, such as setting the **Admin Status** to **Up** or choosing the **Data VLAN** or **Voice VLAN**.

Port reservation editable settings are as follows:

- **Reserve Port**—Enables the port reservation task for the new object;
- **Device Name**—Shows the name of the chosen devices, which must be selected by clicking **Select Device** and using the Device Selector window (for information, see [Using the Device Selector](#));
- **Interface**—Drop-down menu listing for all interfaces on the selected device.  
**The Following VLANs Are Configured**—A read-only panel that shows the VLANs, if any, that are configured on the currently selected **Interface** setting.

Port configuration and VLAN settings cannot be performed when editing objects—you are limited to selecting a different port (from the same device or from a different device) to be bound to the current object.

#### Defining Port Reservations for an Infoblox Grid Member



#### Note

Editing Grid members does not allow for port configuration when you create or change a port reservation. For Grid member editing, you can select interfaces but you cannot edit them, such as setting the **Admin Status** to **Up** or choosing the **Data VLAN** or **Voice VLAN**. This section applies only to creating and defining port reservations and configurations for new Grid members. For the complete Grid member creation procedure, see [Adding Grid Members](#).

You can configure port reservations for a Grid member, including HA members, in the Add Grid Member wizard. All interfaces on a member (LAN1, LAN2, HA and MGMT) may have independently defined port reservations. For Grid members, port reservations are not subject to scheduling and workflow approval, and Grid Manager executes them immediately.

### Defining Port Reservations for an HA Pair

The process of defining port reservations and port configurations for a Grid member allows these settings to be defined during the creation of the new member. (You can also edit them afterwards.) Before performing this procedure, consult the sections [Planning for an HA Pair](#), [About HA Failover](#) and [Adding an HA Member](#) for more details.

1. Go to **Grid** → **Grid Manager** and choose **Add** → **Grid Member** from the vertical toolbar. Define your new HA Pair's settings and click **Next**.
2. In the second Add Grid Member Wizard step (Step 2 of 5), click the **High Availability Pair** option and enter the required **Virtual Router ID** value.
3. As with a normal HA pair configuration, enter the IP information about the following interfaces: VIP, Node 1 HA and LAN1 ports, Node 2 HA and LAN1 ports. The VIP address and the IP addresses for all the ports must be in the same subnet. Follow the guidelines provided in the section [Adding an HA Member](#).
4. Ensure all settings are correct and click **Next**.
5. The following step in the Add Grid Member Wizard presents a workflow that must be performed three times for each appliance—once for each of the three interfaces (LAN1, HA and MGMT) participating in the HA pair. (LAN2 may be configured by editing the Grid member afterwards.)

## Beginning the HA Pair Configuration's Port Reservations

6. Begin by clicking the checkbox for the **Node 1 → LAN1** appliance port. (You do not perform port reservation for the VIP port).
  - a. Select the **Reserve Port** checkbox. The **Node1→ LAN1** port listing changes to read **Pending**.
  - b. Click **Select Device** to choose the device (switch or switch-router) that is Node 1 of the HA Pair. For information, see [Using the Device Selector](#).
  - c. After selecting the device, click **OK**.
  - d. Choose the device port for Node 1's LAN1 port by choosing it from the **Interface** menu. The panel **The Following VLANs are configured** refreshes to show any VLANs that are currently provisioned on the selected port.
  - e. (*This step is optional.*) Select the **ConfigurePort** checkbox.
  - f. (*This step is optional.*) If the **Data VLAN** setting is enabled, select the **VLAN** from the **Data VLAN** menu.
  - g. (*This step is optional.*) Choose an **AdminStatus** of **Up**.
  - h. (*This step is optional.*) Enter a **Description**.
    - i. For the **Node 1 → HA** port, follow steps 6a through 6h for that appliance port. Ensure you select the same switch and **Data VLAN** settings.
    - j. For the **Node 1 → MGMT** port, follow steps 6a through 6h for that appliance port. Ensure you select the same switch and **Data VLAN** settings.
7. For the second appliance Node 2, follow steps 6a through 6j above.
8. Click **Next** when finished with the Reserve Port and optional Configure Port settings for all interfaces. The result appears similar to the example shown in the below figure.

### Port Reservations for an HA Pair

Add Grid Member > Step 3 of 5

**Node 1**

NETWORK INTERF...	RESERVED INTE...
LAN1	Gi1/0/1
HA	Fa0/1
MGMT	Fa1/0/11

**Node 2**

NETWORK INTERF...	RESERVED INTE...
LAN1	Fa0/12
HA	Gi1/0/13
<input checked="" type="checkbox"/> MGMT	Fa0/15

**MGMT**

Reserve Port Reserving a Switch Port does not ensure its availability

Device \*Name: QAlab2sw02.net.com Select Device Clear

\*Interface: Fa0/15

The following VLANs are configured:

ID	Name
20	to_QAlab2sw02

Configure Port

Data VLAN: 805 QA\_test805

\*Admin Status: Up

Description: NODE2 MGMT connection

9. In the final step of the Add Grid Member Wizard, *if you have defined* port configuration, a final wizard step shows that the port configuration (which is a Port Control task) executes Now, without allowing a scheduling of the task. *Port reservation for the HA Pair is not scheduled* and also takes effect automatically as a part of the new Grid member configuration. (Grid Manager creates the Grid member immediately after you click **Save & Close**.)
10. After you click **Save & Close** to create the new HA pair Grid member, thus closing the Add Grid Member wizard, opening the editor for the HA member and clicking the **Port Reservations** tab displays the LAN2 port along with the previous port settings:

#### Completing the HA Pair Port Reservations with the LAN2 Ports

probe.com (Grid Member Properties Editor)

- Toggle Basic Mode
- General
- Licenses
- Network
- Anycast
- Security
- DNS Resolver
- Monitoring
- System Backup
- SNMP
- SNMP Threshold
- Notifications
- Email
- Pre-Provisioning
- Port Reservation
- Extensible Attributes
- Permissions

**Basic**

Reserve Port Reserving a Switch Port does not ensure its availability.

\*Device Name: Select Device Clear

\*Interface: Select Available To configure this interface, go the Devices tab and open the Interfaces tab for the selected device.

The following VLANs are configured:

ID	Name
----	------

Cancel
Save & Close

11. Select the LAN2 port for each node and follow steps 6a through 6h above. If VLANs are provisioned on the selected port, you can also select them.
12. Click **Save & Close**.



## Editing Port Reservation Settings for an Infoblox Grid Member

Editing an existing Grid member's port reservation does not permit configuring the currently selected port. You can select ports from any device to change the port reservation. For any selected device, ports with an Operational Status of Down appear in the **Interface** list.

For editing Grid member settings, you can select interfaces but you cannot configure them; setting the **Admin Status** to **Up** or choosing the **Data VLAN** and changing the **Description** are not allowed when editing a Grid member.

### Editing a Grid Member's Port Reservation Settings

The screenshot shows the 'Grid Member Properties Editor' window for 'probe-.com'. The 'Basic' tab is active. A 'Reserve Port' checkbox is present, with a yellow warning message: 'Reserving a Switch Port does not ensure its availability.' Below this, there is a '\*Device Name' field with 'Select Device' and 'Clear' buttons. An '\*Interface' dropdown menu is set to 'Select Available', with a yellow warning message: 'To configure this interface, go the Devices tab and open the Interfaces tab for the selected device.' A table titled 'The following VLANs are configured:' has columns for 'ID' and 'Name'. A 'Cancel' button is at the bottom left.

You can edit any appliance's port reservation settings:

1. Go to **Grid** → **Grid Manager** and select the checkbox for the Grid member you wish to edit. Click **Edit** from the vertical toolbar.
2. Click the **Port Reservation** tab.
3. Select the **Reserve Port** checkbox. The **Node1->LAN1** port listing changes to read **Pending**.
4. Click **Select Device** to choose the device. (For information, see [Using the Device Selector](#).)
5. After selecting the device, click **OK**.
6. Choose the device interface from the **Interface** menu. Click **Save&Close** when finished.

On HA pairs the process is quite similar except that you may edit up to the full complement of interfaces used for each appliance in the HA pair. The Grid Member Properties Editor supports all ports for each appliance in the HA Pair.



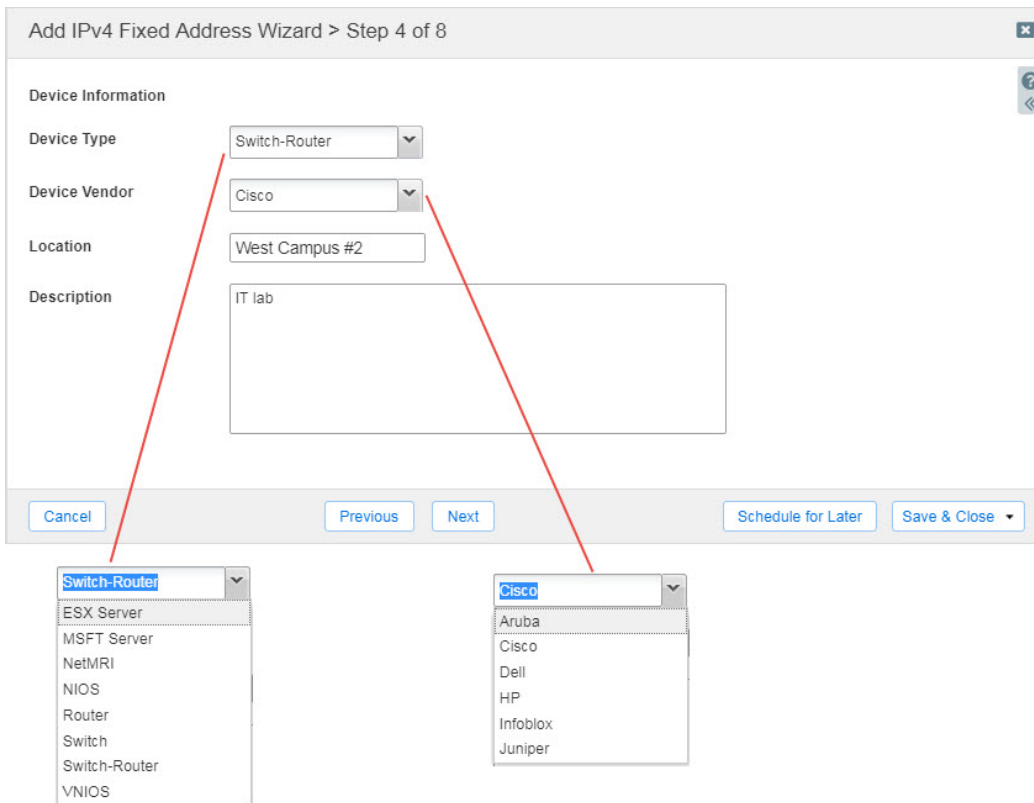
#### Note

When editing Grid members and HA Grid members, Grid Manager does not allow changes to VLAN settings, Admin Status or port description in the editor. You can change these values in the **Interfaces** table by double-clicking the table row for the interface you want to edit.

## Defining Device Information

During the process of configuring a port reservation for an IPAM object, you can define device information settings as descriptive information for IPAM objects, as seen in the below figure.

### Device Information Settings for IPAM Objects in the Wizard



**Note:** If you define Device Information and Device Vendor settings for a port reservation, which are optional, and choose to discover the object to which the port reservation is associated, you may see a conflict if discovery of the object finds that the device for the port reservation is a different type or different vendor. For example, if the specified information states that the device is a switch, and the discovered device is a router, Network Insight reports a conflict. Another example: you declare the vendor to be Aruba, and the discovered value is Cisco. For information on resolving such conflicts, see [Resolving Port Reservation Conflicts](#).

The **Device Type** menu provides the following settings:

- **ESX Server**—identifies the object as a VMware ESX host.
- **MSFT Server**—identifies the object as a Microsoft Hyper-V host.
- **Net MRI**—NetMRI appliance from Infoblox;
- **NIOS**—Infoblox appliance;
- **Router**—Provides routed connections, including VLANs;
- **Switch**—provides L2 switched connections only;
- **Switch-Router**—Provides L2/L3 switched connectivity, including VLANs;
- **VNIOS**—Infoblox virtual appliance.

The **Device Vendor** menu provides the following device vendor choices:

- Cisco
- Juniper
- Infoblox
- Aruba
- Dell
- HP

IPAM object wizards allow definition of the settings when you create the object, or you can define the settings later.

## Editing Interfaces in a Device



### Note

Voice VLAN settings are applicable only for Cisco devices.

To speed port configuration workflows, you can select one interface or multiple interfaces for a device to change the admin status, description and VLAN settings. For example, this feature is handy if you want multiple interfaces to participate in the same data VLAN. There are two ways to approach this feature: directly from the **Devices** page, or by selecting a device on the **Devices** page, opening its **Interfaces** page and selecting ports from there. Editing interfaces is done from the main **Data Management** → **Devices** page.

1. From the **Data Management** tab, select the **Devices** tab.
2. Click the Action icon

for a chosen device and choose **Interfaces** from the popup menu.

3. Select the checkbox for a specific interface, click the Action icon

for the interface and choose **Edit**. The interface editor appears as shown in the figure below.  
*Interface Editor, with editable Admin Status and Description settings*

WS-C2960S-24TS-L.inca.infoblox.com > Gi1/0/5 (Interface)

Basic

General  
VLANs  
Extensible Attributes

IP Address

MAC Address E8:ED:F3:A3:2A:85

Admin Status Up

Operation Status Down

Description test-12

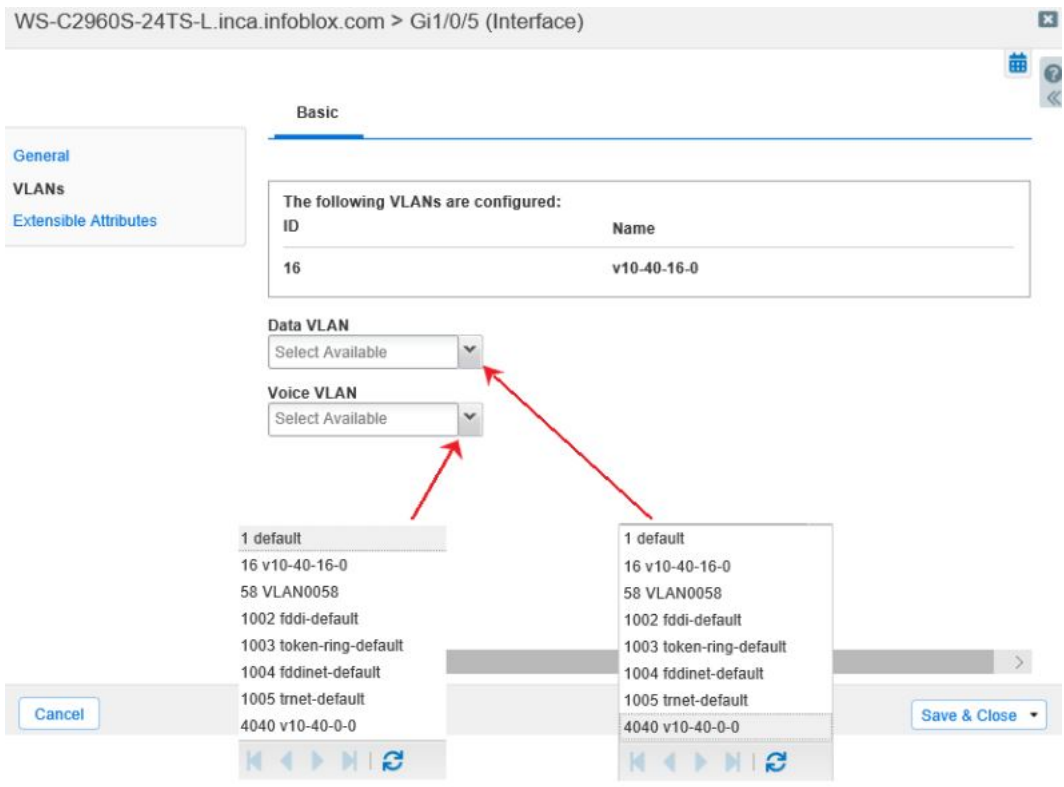
Port Type: ethernet-csmacd  
Port Speed: 10 Mbps  
Trunk Status: Off  
Link Aggregation: No  
NIOs Management Status:  
IPAM Type:  
Usage:

Cancel Save & Close

4. The editable settings are **Admin Status (Up or Down)** and **Description** (click inside the field to edit). In some cases, owing to device permissions, the device type or other device settings, you may not be able to edit these values for the selected interface.

5. To edit VLANs for the chosen interface, click the **VLAN** tab. The following figure shows an example. VLAN editing also is subject to permission limitations based on the device, and on the device type.

### Editing VLANs



6. Choose a data

VLAN to assign to the port from the **Data VLAN** drop-down menu.

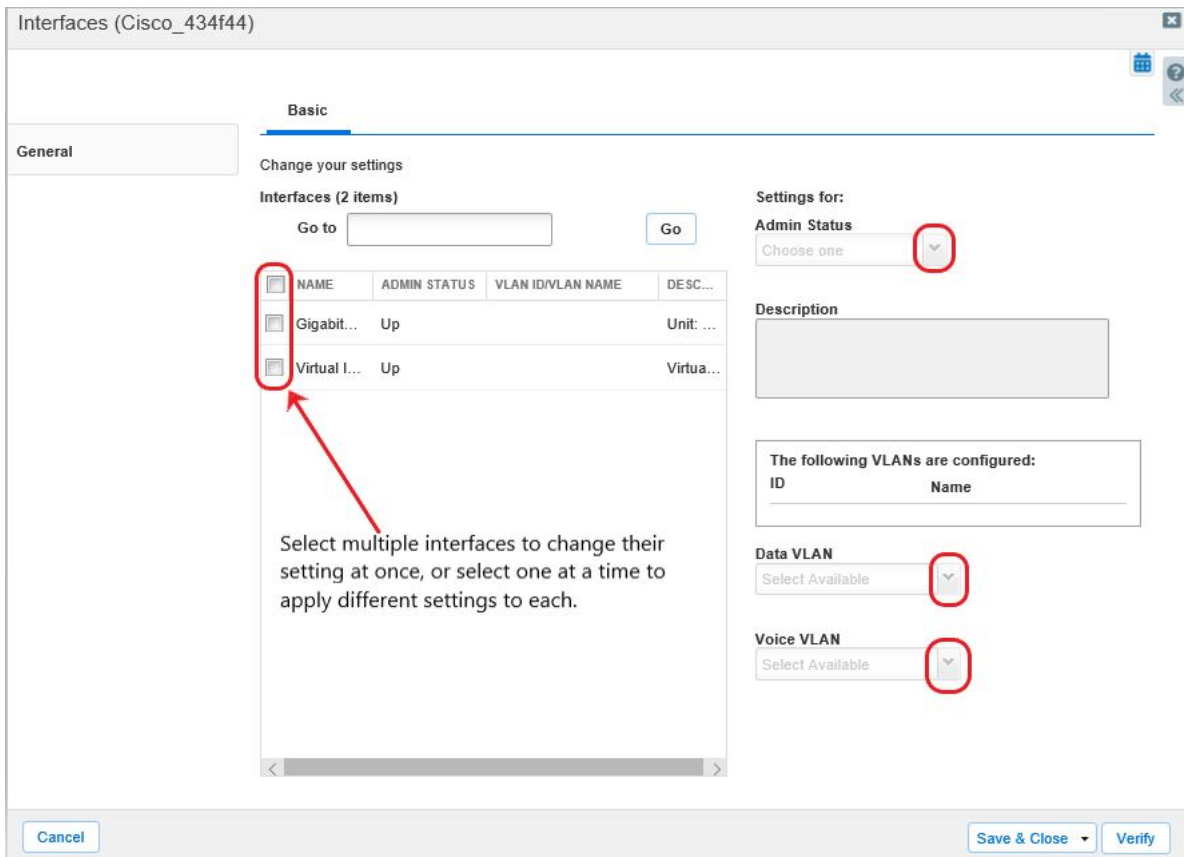
7. If supported, choose a voice VLAN (Cisco only) from the **Voice VLAN** drop-down menu.
8. Click the **Extensible Attributes** tab to add any attributes that are necessary for the interface.
9. Click **Save & Close** to close the interface editor.

## Editing Multiple Interfaces

To edit multiple interfaces for a device:

1. From the **Data Management** tab, select the **Devices** tab.
2. Click the Action icon
  - for a chosen device and choose **Interfaces** from the popup menu.
3. Select the checkboxes for each interface that you want to edit.
4. Expand the Toolbar and click **Edit**. The Interfaces editor appears as shown in the figure below.
 

*Editing Multiple Interfaces*



You can select one or more interfaces for configuration, define their settings in the dialog, and then select other interfaces and define different settings for them.



#### Note

Once you change **Admin Status**, **Data VLAN** or **Voice VLAN** settings for any selected port, no automatic reversion exists to the original settings from the same editing session. You must cancel out of the Interfaces editor to reject any changes and begin with a new editing session from the **Interfaces** page. Use the **Verify** button to verify your changes.

5. Select the checkboxes for one or more ports and define the Port Configuration settings for the following:
  - **Admin Status:** Select Up or Down from the menu, depending on the current state of the port(s);
  - **Description:** Provide a brief description of the port configuration or other information;
  - **DataVLAN:** (Hidden if editing a VLAN is not supported) Drop-down list of all data VLANs actively configured in the current device. One of the values can be chosen for the currently select interface(s);
  - **VoiceVLAN:** (Hidden if editing a voice VLAN is not supported) Drop-down list of all voice VLANs actively configured in the current device. One of the values can be chosen for the currently select interface(s).
6. After making configuration changes to all ports, click **Verify** to check over your changes:

## Verify your new settings

**Changed settings**

NAME	ADMIN STATUS	VLAN ID/VLAN NAME	DESCRIPTION
G03/26	Up	9 vlan-009	GigabitEthernet3/26
G03/42	Up	21 ICIKA	GigabitEthernet3/42
G03/43	Up	22 FGHGH	GigabitEthernet3/43

7. Click **OK**. The changes are not committed by doing so.

8. If the port configuration changes are correct, click **Save & Close** or click the Scheduling icon at the top of the editor. To schedule this task, click the Schedule icon at the top of the editor. In the *ScheduleChange* panel, click **Later**, and then specify a date, time, and time zone. The Schedule icon is green when there is a pending scheduled task. You can reschedule the task if you have the applicable permissions.

When you complete the configuration, all port configurations in the session are combined into a single task by Grid Manager.

### Editing Multiple Interfaces from the Devices Page

The **Data Management** → **Devices** page allows you to immediately set the Admin Status, descriptions and VLAN settings for any selection of ports on a single device.

1. Click the Action icon

for a chosen device and choose **Edit** from the popup menu. The **Interfaces** page appears for the device editor.

2. Select the checkboxes for one or more ports and define the Port Configuration settings for the following:
  - a. **Admin Status:** Select Up or Down from the menu, depending on the current state of the port(s);
  - b. **Description:** Provide a brief description of the port configuration or other information;
  - c. **Data VLAN:** (Hidden if editing a VLAN is not supported) Drop-down list of all data VLANs actively configured in the current device. One of the values can be chosen for the currently select interface(s);
  - d. **Voice VLAN:** (Hidden if editing a voice VLAN is not supported) Drop-down list of all voice VLANs actively configured in the current device. One of the values can be chosen for the currently select interface(s).
3. After making configuration changes to all ports, click **Verify** to check over your changes.
4. Click **OK**. The changes are not committed by doing so.
5. If the port configuration changes are correct, click **Save & Close** or click the Scheduling icon at the top of the editor. To schedule this task, click the Schedule icon at the top of the editor. In the *Schedule Change* panel, click **Later**, and then specify a date, time, and time zone. The Schedule icon is green when there is a pending scheduled task. You can reschedule the task if you have the applicable permissions.

When you complete the configuration, Network Insight combines all port configurations in the session into a single task.

### Inline Interface Editing

You can directly edit a device's interface listings in the Interfaces page. Double-click any table row and the fields that you can edit for the row, which consist of the VLAN ID, Admin Status and **Description** fields for each interface, open in an inline editing selection field.

Fields are editable as applicable to device types. If you are editing an interface on a L2 switch, for example, the **VLAN ID** field does not appear.

### Editing Interface values inline.

NAME	IP ADDRESS	VRF...	NETWORK VIEW	V...	MAC ADDRESS	ACTIVE USERS	VLAN ID	VLAN NAME	PORT TYPE	PORT SPEED	ADMIN STATUS	OPERATION STATUS	TRUNK...	DESCRIPTION
StackSub-St1-2			default			0			propVirtual	0 bps	Up	Down	Off	StackSub-St1-2
StackPort1			default			0			propVirtual	0 bps	Up	Down	Off	StackPort1
Nu0			default			0			other	10 Gbps	Up	Up	Off	Null0
StackSub-St1-1			default			0			propVirtual	0 bps	Up	Down	Off	StackSub-St1-1
Gi1/0/1			default		E8:ED:F3:A3:2A:81	0			ethernet...	10 Mbps	Up	Down	Off	GigabitEthernet1/0/1
Gi1/0/2			default		E8:ED:F3:A3:2A:82	0	1	default	ethernet...	10 Mbps	Up	Down	Off	GigabitEthernet1/0/2
Gi1/0/3			default		E8:ED:F3:A3:2A:83	0	58	VLAN0058	ethernet...	10 Mbps	Up	Down	Off	FA-test
Gi1/0/4			default		E8:ED:F3:A3:2A:84	0	1	default	ethernet...	10 Mbps	Up	Down	Off	GigabitEthernet1/0/4
Gi1/0/5			default		E8:ED:F3:A3:2A:85	0	16	v10-40	ethernet...	10 Mbps	Up	Down	Off	test-12
Gi1/0/6			default		E8:ED:F3:A3:2A:86	0			ethernet...	10 Mbps	Up	Down	Off	Stanley GigabitEther...

Double-clicking a table row opens the editable fields for the selected record. If editable fields are not present in the table display, you cannot change their values in the **Interfaces** page. After making inline changes, click **Save** on the selected row to commit them. To prevent using any changes, click **Cancel**. This also de-selects the row.

|| jsr111

Task state: Pending

Current value: jsr111

Requested value: jsr111\_testonly

[Go to task manager](#)

**Note**

When you make inline changes to an interface, a new task is created under Grid Manager, which you can view in the Task Manager page (for more information, see [Viewing Tasks](#)). A status icon appears next to the interface element you have changed, indicating the status of the new task and providing a link to the Task Manager page. New tasks appear with a status icon of Pending (||). When the new task completes, the icon changes to a green checkmark.

## Monitoring Device Lifecycle and Vulnerabilities Using Advisor

Advisor is a subscription-based service that assists you in monitoring and maintaining network and security infrastructure based on released Common Vulnerabilities and Exposures (CVEs) and vendor product lifecycle announcements. It allows you to monitor the following equipment lifecycle and vulnerability data for the devices discovered by Network Insight:

- CVE Bulletin Count: The count of CVEs that affect the device.
- CVE Bulletin List: The names of manufacturer's CVEs that affect the device.
- End of Sale Date: The date announced by the manufacturer when the device cannot be ordered.
- End of Support Date: The date announced by the manufacturer when the device is no longer eligible for support.
- EOL Bulletin URL: The URL pointing to the manufacturer's announcement.
- EOL Status: The End of Life status of the device, e.g. Current, End of Sale, and End of Support.

You can execute Advisor at any time or configure it to run at specific intervals and time. For information, see [Configuring Advisor Properties](#). In addition, you can use the "Device Advisor" report and [dashboard](#) and Advisor filters in Smart Folders.

You can view lifecycle and vulnerability data in the **Devices** tab. If not present, add the Advisor-related columns to the table. To do so, click a column header, click the down arrow, select **Columns** -> **Edit Columns**, and then select the Advisor-related columns. You can also view this data in the properties of any discovered device as described in the next section.

Also see [Advisor Vendor Support Matrix](#) below.

## Viewing Advisor Data

To see lifecycle and vulnerability data for an individual device:

1. In the Grid Manager, select **Data Management** → **Devices**.
2. Select a device in the table and click **Edit**.
3. In the **Device Properties** editor, select **Extensible Attributes**. The **Value** column displays the lifecycle and vulnerability data, if detected for the device.



### Note

Unlike regular extensible attributes that you can create to track NIOS objects, Grid Manager creates the Advisor extensible attributes automatically to obtain and display device lifecycle and vulnerability data.

## Advisor Vendor Support Matrix

The following table lists vendors that support Advisor. It includes supported announcement types and the NIOS version when they were added or updated.

Vendor	Vulnerabilities	Lifecycle
<b>8.4</b>		
Cisco	Yes	Yes
<b>8.4.4</b>		
Arista	Yes	Yes
Blue Coat	Yes	Yes
Checkpoint	Yes	No
Infoblox	Yes	Yes
Juniper	Yes	Yes
Palo Alto Networks	Yes	Yes
<b>8.6</b>		
Aruba	Yes	Yes



Vendor	Vulnerabilities	Lifecycle
Avaya/Nortel	Yes	Yes
Brocade	Yes	Yes
Checkpoint (API updated)	Yes	Yes
Netscreen (same API as Juniper)	Yes	Yes

## Viewing Discovery Status



### Note

Opening Discovery Status for viewing requires Superuser permissions under Grid Manager.

You can view the complete discovery status of all devices or of selected devices.

To isolate devices for evaluation, use filtering to reduce the list. Click **Use Filter** at the top of the table and choose **IP Address**, **Name**, or **Overall Status** as the filter.

To view the discovery status, complete the following:

1. From the **Data Management** tab, select the **IPAM** or **Devices** tab.
2. In the Toolbar, click **Discovery Status**. The Discovery Status table is displayed.

The Discovery Status table lists detailed information about network devices and end hosts discovered through all methods, including SNMP, ICMP ping sweeps, and other processes. It includes:

- **IP Address:** the IPv4 or IPv6 address of the discovered device. You can filter the table by this value.
- **Name:** The name of the discovered device as reported through SNMP. You can filter the table by this value.
- **Type:** The discovered device type. Examples include **Router**, **NIOS**, **Switch-Router**, **Firewall**, **Load Balancer**, **LWAP**, and numerous others.
- **Overall Status:** Indicates the overall success or failure of the discovery task on the device. Hover the mouse over the device to see more detailed information about the discovery status, including the timestamp of the last discovery event, confirmation of detection ("Device Exists"), and the means of detection, which are usually methods such as SNMP, reading the ARP table or location through a Seed router. You can filter the table by this value.
- **Reached Status:** Indicates the reachability of the discovered device. Typically, devices are reported **Passed** for **Reached Status** if they are reachable through SNMP, a path trace through ICMP, or UDP-based path tracing for an IPv6 address.  
You may see a **Reached Status** of **Passed** and still receive an **Overall Status** of **Failed**. This often occurs because either the CLI credentials or SNMP credentials provided for discovering the device do not work, or another problem occurs in some part of the discovery process.
- **SNMP Collection Enabled:** Indicates whether the managed device allows SNMP as a management protocol. This value shows **Yes** or **No**. You do not see any SNMP collection status updates if this value shows No.
- **SNMP Credential Status:** Indicates whether the correct SNMP credential is used by discovery. Usually shows simple **Passed** or **Failed** status. Passing the mouse over the Failed status reading for this column shows the location in the SNMP data collection where information gathering failed. The typical message for a failure of this type shows Failure to Authenticate, which simply means that the correct SNMP credentials have not been provided for either SNMPv1, SNMPv2c, or SNMPv3 as required and defined for the device's discovery configuration. The tooltip also displays the name of the credential group that was used to guess credentials for the device.  
If **SNMP Credential Status** shows Failed, you do not see a value under **SNMP Collection Status**, because that is

dependent on successful credential authentication. Should you succeed in SNMP Credential Authentication for a device, this value shows **Passed**.

- **SNMP Collection Status:** Indicates whether managed device information has been successfully collected from the device. If the current device shows an **SNMP Credential Status** of Failed, this field remains blank. Should you succeed in SNMP Credential Authentication, **SNMP Collection Status** may or may not show a Passed outcome. If the final outcome is successful, passing the mouse over the table value shows the SNMP data set that was successfully collected from the device. When the **SNMP Collection Status** is set to **Passed**, it indicates that the credentials provided in *Member Discovery Properties* for discovering the device are correct and discovery has been completed successfully on the device. For more information, see [Defining the Discovery Member Type](#).
- **CLI Credential Status:** Reports the basic success state of CLI credential usage for device discovery. When you see a Failed status in this column, hover the mouse over the table value. Details related to failed CLI credentials normally relate to "Failed to Authenticate" events. (For more information, see the following subsection Analyzing Discovery Status.) If you define device discovery requirements to use both SNMP and CLI, and you receive complete SNMP discovery information but fail to authenticate for CLI, the **Overall Status** for the device remains as **Failed**. The tooltip also displays the name of the credential group that was used to guess credentials for the device.
- **CLI Collection Enabled:** indicates the CLI collection configuration state for the discovered or managed device. Possible values are **Yes** (CLI collection is enabled) or **No** (CLI collection is disabled for the device).
- **Fingerprint Status:** Shows the status of discovery of the device's OS through fingerprinting.
- **Last Update:** Timestamp showing the conclusion of the last data update for the current device.
- **First Seen:** Timestamp showing the initial discovery event.
- **Last Seen:** The date and time when the device was last successfully polled by discovery.
- **Last Action:** The last action performed by discovery upon the device after the discovery took place. Hover the mouse over this field to obtain details.

Visible columns can be changed in the Discovery Status window. At the top of any column header, click the down arrow tool, and choose **Columns** → **Edit Columns**.

## Analyzing Discovery Status

When you see the Failed status for a device under **Overall Status**, the problem usually relates to issues in discovery data collection. When **Overall Status** shows a value of **Passed**, it indicates that everything has passed for the device. If the value is Failed, it indicates that one or more elements of the device have failed discovery. Your discovery settings have a great deal to do with what you see in the respective Status columns. You key on the Overall Status result and read columns to the right to narrow down possible causes. To start, you typically see three basic discovery credential configurations for network devices:

- SNMP credentials and no CLI credentials;
- SNMP credentials and CLI credentials;
- No SNMP credentials, CLI credentials only.

Begin by considering SNMP-only device discovery configurations. In the table, **Overall Status** shows only a simple Failed message with no detail. Go to the next data column, which is **Reached Status**. If the device proves reachable, this value shows Passed, indicating that discovery can successfully reach and query the device. If **Reached Status** shows Failed, this is the first and most fundamental problem, that the device cannot be Pinged or contacted in any way across the network.

As an example, assume a **Reached Status** value of Passed. If the device is reachable, discovery can successfully attempt SNMP or CLI communication to the device. Beginning with SNMP, and assuming that SNMP collection is enabled, select the **SNMP Credential Status** counter. If it shows Failed, that normally indicates a Failure to Authenticate, which can be shown as a tooltip by hovering the mouse over the table field:

2014-04-21 09:48:53 PDT : SNMP Credentials: Failed to authenticate

2014-04-21 10:46:51 PDT : SNMP Credentials: Successfully authenticated / Version: SNMPv2c

A successful SNMP Authentication may or may not result in successful data collection:

2014-04-21 09:47:52 PDT : SNMP Collection: Failed to collect data / Table: Forwarding

2014-04-21 14:17:00 PDT : SNMP Collection: Failed to collect data / Table: System

A successful authentication also shows which protocol was used for SNMP authentication; SNMPv1, SNMPv2c or SNMPv3. This does not guarantee successful data collection through SNMP, however.

The **SNMP Collection Status** counter shows possible Passed and Failed values. Hovering the mouse over a Failed value in this column shows a tooltip reporting the set of data that discovery could not collect through SNMP. When discovery encounters an error during collection of a specific data set (Forwarding table, or System identification data, for example), data collection stops and issues an error message and an SNMP trap, which is reported and also appears in the tooltip. If **SNMP Collection Status** shows a value of Passed, and discovery does not use CLI data collection on the device, discovery has successfully completed on the device.

The **CLI Credential Status** counter also reports either Passed or Failed results, and uses tooltips to tell the user what is going on in more detail.

2014-04-21 09:56:50 PDT : CLI Credentials: Failed to authenticate / Connection closed by foreign host

2014-04-21 10:00:27 PDT : CLI Credentials: Failed to authenticate / Username and password ok, but timed out during prompt detection

2014-04-21 09:55:21 PDT : CLI Credentials: Failed to authenticate / Connection refused

CLI credentials failure messages are straightforward and can be tested by verifying login tuples or Enable passwords from the credential sets defined in discovery configuration.

If you receive a **CLI Credential Status** value of Passed, the correct command-line admin login information is specified in the discovery configuration.

For more information about checking and diagnosing discovery behavior for devices listed in the status table, see the following topic [Executing Discovery Diagnostics](#).

## Provisioning and De-Provisioning Networks

You may provision networks and remove or de-provision networks from individual devices. You can also provision networks when creating a network in IPAM (for more information, see [Adding IPv4 and IPv6 Network Containers and Networks](#)). Network provisioning and de-provisioning comprises the second type of port control tasks under Network Insight. Provisioning a network involves creating a new network and adding it to the list of networks in IPAM, and also involves changes to device configuration.




#### Note


If a port control task requires administrative approval, and it is not approved before its scheduled execution, the task appears as unsuccessful in Task Manager.

Provisioning networks also allows for provisioning VLANs. If devices do not support VLANs, options for provisioning VLANs do not appear for those devices and their associated interfaces.

If a network is already provisioned for an interface, regardless of its status under Grid Manager, you cannot provision another network upon it.

Only available interfaces that support network provisioning are shown in Grid Manager for provisioning tasks. The horizontal toolbar also provides related functions:

**Provision Network** : Available for discovered devices and for managed devices, this icon opens the **Provision Network** feature, allowing you to provision an existing IPAM network onto the selected device by selecting a device interface or assigning a VLAN. Grid Manager creates a new Port Control task, and you can choose the interface on which the network is provisioned, along with VLAN configuration and other settings.

**De-provision Network** : Available for networks that are managed under IPAM, for de-provisioning on devices that are managed under IPAM on the **Data Management** → **Devices** page. A dialog box appears summarizing the task you are instructing Grid Manager to perform. This action changes the configuration of the device.

To provision IPAM networks onto a device:

1. From the **Data Management** tab, select the **Devices** tab.
2. Click the **Next Page** and **Last Page** icons to locate the device whose interfaces you want to provision.
3. Click the Name of the device. The Devices page displays the five tabs of information associated with the device.
4. Click the **Networks** tab.
5. The Networks page lists all discovered networks (highlighted in grey), unmanaged networks (highlighted in yellow), and any managed networks (highlighted in white, showing **Yes** in the **Managed** column) present on the selected device.
6. Open the vertical toolbar and click **Provision Network**. The Provision network wizard appears.
7. Click **Select Network** to choose the network the you want to provision. If only one managed network is available on the device, that network value is populated after clicking the button.  
If more than one managed network is available under IPAM, the Network Selector dialog opens, listing all networks managed under IPAM.
8. Choose the network to provision onto the device and click **OK**.
9. Enter the **Router IP Address**. This required field may be pre-populated with the DHCP router IP address if the device already has a DHCP configuration. If not, enter the gateway router IP address for the current device.
10. If necessary, check the **DHCP Forwarding** checkbox. Check this checkbox to enable DHCP forwarding for the newly provisioned network. If a DHCP failover is already present, the IP addresses from that failover are used for DHCP forwarding information.
11. For choosing the Interface, you can choose one out of two options:
  - a. **Interface** drop down list: If you are provisioning the network directly onto an interface, select it from this list. Only interfaces that are available for provisioning on the chosen device appears on this list; interfaces that are already active in a network do not appear;  
–or–
  - b. For a switch-router, select **Create VLAN**, and specify the **VLAN Name** and its new **VLAN ID**. Ensure that the VLAN ID is not one that is already provisioned on the device.

*Provisioning a VLAN on a Switch-Router*

Provision Network > Step 1 of 2

\*Network

Device WS-C3750X-24.inca.infoblox.com

\*Router IP Address

DHCP Forwarding  DHCP Forwarding will be used only when applicable.

Interface

Create VLAN

VLAN Name  VLAN ID

12. Click **Next** to go to the second step in the Provision Network wizard, in which you define whether to provision the configuration now or to schedule it.
13. To immediately provision the new network on the chosen device, select **Now**.
  - a. You can choose to have Grid Manager create the network at a later time. To do so, select **Later**. Choose a **Selected** time by entering or selecting a **Start Date** (click the calendar icon to choose a calendar date) and a **Start Time**, and choose a **Time Zone**.
14. Click **Save & Close** when finished.

## De-Provisioning Networks



### Note

De-provisioning a network changes the device configuration. As such, a separate task is created for the action under Task Manager. However, you cannot schedule the de-provisioning of a network—once you confirm the de-provisioning action in Grid Manager, the action takes place. Each managed and unmanaged device under Grid Manager provides a Permissions page (**DataManagement**→**Devices**→ Select Device → click **Edit**→**Permissions** tab). By default, no admin group or Role is assigned to managed devices. Infoblox recommends using caution when assigning rights to users that may be able to access devices and change device configurations.

De-provisioning networks is a relatively straightforward task that can be performed for any selected network, whether it is a non-NIOS network (a network that cannot be configured in IPAM), an unmanaged network, or a managed network.



### Note

If the network is also managed under IPAM, de-provisioning the network from a device does not delete the network from IPAM.

If you are deleting a network from the main IPAM page, any devices that have endpoints provisioned on that network are also de-provisioned for that network.




#### Note

A network may not be de-provisioned until after you set the interface for the network on the device(s), to **Down** in **Admin Status**.

1. From the **Data Management** tab, select the **Devices** tab.
2. Click the **Next Page** and **Last Page** icons to locate the device through which you want to locate the interfaces to convert.
3. Click the Name of the device. The Devices page displays the five tabs of information associated with the device.
4. Click the **Networks** tab for the chosen device. The Network page lists all discovered networks (highlighted in grey), unmanaged networks (highlighted in yellow), and managed networks (highlighted in white) present on the selected device.
5. Click the **Next Page** and **Last Page** icons to locate the network that you wish to de-provision.
6. Click the Action icon

next to the network you want (this automatically selects it) and select **De-Provision Network**. The dialog box appears, listing the device name, the device's IP address, the interface to which the network is currently bound, and the network's endpoint IP address on the current device.

#### *De-Provisioning a Network from a Device*



7. Click **Yes** to confirm the de-provisioning action.



#### Note

Ensure that the de-provisioning of the network has administrative approval.

You can also select multiple network entries from the list on the same device and de-provision all of them in a single step. Exercise caution when performing such actions.

### De-Provisioning Networks by Deleting Networks in IPAM



#### Note

Deleting a network under IPAM creates a new Object Change task in Task Manager. You can check the **Administration→Workflow→TaskManager** page to view its status.

You can simply delete a managed or unmanaged network in IPAM to de-provision it. Doing so opens a Delete Confirmation dialog. IPAM also automatically prompts you to verify that you are deleting the network from all devices that have interfaces connecting to the network, subject to verification and permissions.

By default, when you delete the network, all devices that connect to the network, *that are also managed by IPAM*, are part of the new de-provisioning port control task created by Grid Manager. If you do not want the network de-provisioned from all devices, clear the **De-provision network from all interfaces** checkbox or simply cancel out from the Delete Confirmation dialog.


### Troubleshooting Port Control Tasks

Issues can occur when attempting to define port configurations on devices. When you define port configurations through Network Insight, you are defining a *port control task* that can be viewed, investigated, and run again when necessary.

You do so by using the Task Manager (**Administration** → **Workflow** → **Task Manager**) and looking for tasks that show a **Type of Port Control**. Each Port Control task provides an Execution Log and the ability to re-run a task that has failed for any reason.

The Task Manager page provides an Action icon ≡ column with a series of menu options for features related to Grid Manager tasks to manage task execution, scheduling and approval. Menu choices change based upon the context and the current state of tasks in the table; features available in the Action menu include the following:

#### Task Manager Action menu



- **Approve:** Enables admins to approve a pending job.
- **Reject:** Enables admins to reject a pending job, immediately cancelling it.
- **Execution Log:** Opens a completed task's execution log window. the Execution Log lists the complete communications sequence sent to a device to perform a port control task.
- **Execute Now:** Force a selected pending task to execute immediately.
- **Re-Execute:** Allows you to re-run the selected task. Combined with the Execution Log, this process can aid in troubleshooting a failed port control task.
- **Reschedule:** Opens the Reschedule window for the selected task. To immediately execute this task, click **Now**. Or, in the *Reschedule* panel, click **Later**, and then specify a date, time, and time zone. You can reschedule the task if you have the applicable permissions. Click **Save** to commit the changes.
- **Delete:** Deletes the pending task.
- **View:** Opens the Task Viewer to the currently selected task. For related information, see [Using the Task Viewer to View Job Logs and Approve Jobs](#).

The Execution Log allows you to see task behavior when it executes. You can check the configuration directly on the device and re-run the job if it has failed, by selecting **Re-execute** from the Action menu on the Task Manager page.

## Supported Discovery Methods

When you perform a discovery, you can choose any or all of the following discovery methods:

- SNMPv1/v2c device polling as described in [SNMP](#) below.
- SNMPv3 device polling as described in [SNMP](#) below.
- CLI device querying as described in [CLI](#) below.
- ICMP Ping Sweep and Smart Subnet Ping Sweep as described in [ICMP](#) below.
- TCP as described in [TCP](#) below.
- NetBIOS as described in [NetBIOS](#) below.

These methods actively scan predefined networks and probe IP addresses. The appliance listens for responses from the IP addresses as proof of activity. The IP discovery scans through the specified network ranges and probes IP addresses (except for the network, broadcast, and multicast address types) in each network, including the /31 and /32 subnets. Note that addresses in the /31 and /32 subnets can be used only as source addresses for point-to-point links and loopbacks. In these cases, no broadcast or network addresses exist in the /31 and /32 subnets, and the appliance can discover source addresses in these subnets.



## SNMP



### Note

Infoblox does not recommend using vendor default SNMP credentials on network devices. Should you need to use vendor defaults for a given device type, you enter those values in the list of SNMP credentials on the Grid Master.

Network Insight supports discovery of devices and networks through SNMPv1/v2c and through SNMPv3 protocols. Discovery acquires information from standard SNMP MIB object IDs (OIDs) to correctly identify and catalogue devices. You enter or import lists of SNMP credentials with which the appliances query devices on the network to perform discovery.

SNMPv1 and SNMPv2c protocols are combined into a set termed **SNMPv1/v2** for discovery. SNMPv1/v2 discovery requires standard read community strings to be stored on the Grid Master.

Accounts using SNMPv3 use a standard suite of authentication and security protocols. If Network Insight uses SNMPv3 to collect data from devices supporting the protocol, you can define specific user credentials with combinations of authentication and protocol support, and the unique keys for each protocol. Network Insight also supports multiple entries for the same username string, enabling checking of similar SNMPv3 credentials that use different authentication and security protocols.

Some devices found by discovery may not have known SNMP credentials or credentials that are entered into the sets of SNMP credentials defined for discovery.



### Note

SNMP Credentials from the Grid or from the Member credential list are always tried in the specified order unless a credential is associated with a host, fixed address or reservation being discovered.

## CLI



### Note

CLI is optional for discovery but is required for all Port Control operations. Discovery can perform CLI data collection to collect information for specific device types. SNMP is required for all device discovery.

Network Insight enables the use of dynamically created and closed Telnet and SSH command-line sessions to log in, query, and configure ports using each device's command-line syntax. Network Insight does so without requiring extensive configuration from the user. You need to provide known admin account login information and any Enable passwords for devices in the networks to be discovered. CLI credentials are required for port reservation and port configuration operations under Grid Manager. You enter CLI credentials under Grid Discovery Properties (**Grid → Grid Manager → click Edit → Grid Discovery Properties**) to be inherited by discovery Probe members, and as necessary for each discovery Probe member. You can also override them for individual IPAM objects (fixed addresses, hosts and IPv4 reservations) and test the CLI credentials against devices for correctness. For more information, see [Testing SNMP and CLI Credentials](#).

## ICMP

Discovery uses different variations of Ping traces to perform higher-performance, brute-force device discovery. ICMP is the last resort when devices do not support SNMP management protocols or an SNMP credential is lacking. The ICMP Smart Ping Sweep option enables brute-force subnet Ping sweeps on IPv4 networks. Subnet ping sweeps are used as a last resort in the discovery process. A subnet ping sweep is performed if Network Insight is unable to identify any network devices in a given subnet. Subnet ping sweeps are performed no more than once per day, and will end the



ping sweep on a given subnet once Network Insight discovers a network device and is able to collect data from it. You can configure the timeout value (**Ping Sweep Timeout**) and the number of attempts (**Ping Sweep Attempts**).



#### Note

Smart subnet ping sweeps are not performed on subnets larger than /22. Ping sweeps of any kind do not apply on IPv6 networks because of the greater scale of network addresses in the IPv6 realm.

Complete Ping Sweep differs from the Smart Subnet ping sweep in the following ways:

- The discovery ping sweep runs only against the specified range.
- The sweep runs regardless of the range size.
- The sweep runs regardless of the number of discovered devices within the specified range.

Discovery also performs automatic Ping traceroutes when needed for path collection. Path collections run without user intervention or configuration.

## TCP

TCP scanning probes each active host on a list of TCP ports using TCP SYN packets. This method detects all active hosts that generate SYN ACK responses to at least one TCP SYN. The discovery can determine the OS on a host by analyzing how the host reacts to the requests on opened and closed ports. It then uses the TCP fingerprints to guess the OS. To obtain a TCP fingerprint, IP discovery provides two scanning techniques, SYN and CONNECT.

When you use the SYN technique, the discovery sends a TCP SYN packet to establish a connection on a TCP port. If the port is open, the host replies with a SYN ACK response. The discovery does not close the port connection.

The CONNECT technique is a three-way TCP handshake. The discovery starts with the same process as the SYN technique by sending the TCP SYN packet. A response containing a RST flag indicates that the port is closed. If the host replies with a SYN ACK response, discovery sends a RST packet to close the connection. If there is no reply, the port is considered filtered. TCP scanning is a deliberate and accurate discovery method, enabling detection of all active hosts on a network provided that there are no firewalls blocking TCP packet exchanges.

You can choose the TCP ports and the TCP scanning technique in the *Grid Discovery Properties* editor. This method returns the following information for each detected host:

- **IP address:** The IP address of the host.
- **MAC address:** The discovery returns the MAC address only if the Probe member running the discovery is on the same discovered network.
- **OS:** This is set to the highest probable OS reported in the response.

To use the TCP discovery method, the TCP port and a specific set of ports between the Probe member and the discovered networks must be unfiltered. The default set of ports is defined by the factory settings.

## TCP Port Scanning

By enabling port scanning, Network Insight probes the list of TCP ports enabled in the Advanced tab, to determine whether they are open. You can control some settings for port scanning behavior, including the choice of a TCP scanning technique.

- **Profile Device:** If enabled, Network Insight attempts to identify the network device based on the response characteristics of its TCP stack, and uses this information to determine the device type. In the absence of SNMP access, the Profile Device function is usually the only way to identify devices that do not support SNMP. If you disable Profile Device, devices accessible via SNMP are still correctly identified; all other devices are assigned a device type of **Unknown**. **Profile Device** is disabled by default for discovery polling.

The Profile Device option uses the editable list of TCP protocol ports from the **Grid Discovery Properties → Polling → Advanced** tab as its profile, and polls each of the ports enabled in that list, using the configured timeout value and the number of polling attempts for each port.

For more information, see [Defining Seed Routers for Probe Members](#).

Should you disable Port Scanning, discovery attempts no port probes other than SNMP on any device.

## NetBIOS

The NetBIOS method queries IP addresses for an existing NetBIOS service. This method detects active hosts by sending NetBIOS queries and listening for NetBIOS replies. It is a fast discovery that focuses on Microsoft hosts or non-Microsoft hosts that run NetBIOS services.

NetBIOS discovery returns the following information for each detected host:

- IP address: The IP address of the host.
- NetBIOS name: This value is set to the name returned in the NetBIOS reply.

To use the NetBIOS discovery method, ports 137 (UDP/TCP) and 139 (UDP/TCP) between the Grid member performing the discovery and the target networks must be unfiltered.

The following table summarizes the supported discovery methods:

Discovery Type	Returned Data	Guideline	Mechanism
Smart IPv4 Subnet Ping Sweep	<ul style="list-style-type: none"> <li>• IP addresses</li> <li>• MAC addresses</li> </ul>	Apply on known subnetworks on which no devices are readily found. Limited to networks of /22 and smaller.	ICMP echo request and reply.
Complete Ping Sweep	<ul style="list-style-type: none"> <li>• IP addresses</li> <li>• MAC addresses</li> </ul>	Last resort for discovery. Use ICMP for a rough and fast discovery. Enables path tracing.	ICMP echo request and reply, ICMP traceroute.
NetBIOS	<ul style="list-style-type: none"> <li>• IP addresses</li> <li>• NetBIOS name</li> </ul>	Use NetBIOS for discovering Microsoft networks or non-Microsoft networks that run some NetBIOS services	NetBIOS query and reply.
TCP	<ul style="list-style-type: none"> <li>• IP addresses</li> <li>• MAC addresses</li> <li>• OS</li> </ul>	Use TCP for an accurate but slow discovery	TCP SYN packet and SYN ACK packet.
Port Scanning/ Profile Device	<ul style="list-style-type: none"> <li>• Open and Closed TCP ports</li> <li>• IP Addresses</li> </ul>	Disabled by default, use for non-SNMP devices.	Scans specified list of TCP ports, using TCP SYN packet.

Discovery Type	Returned Data	Guideline	Mechanism
SNMPv1/v2 SNMPv3	<ul style="list-style-type: none"> <li>• Open and Closed TCP ports</li> <li>• IP Addresses</li> <li>• System Description</li> <li>• System Up Time</li> <li>• Routing Neighbors</li> <li>• Routing and Forwarding Tables</li> <li>• ARP tables</li> <li>• SNMP credentials</li> </ul>	<p>Most important protocols for discovery. Ensure you have the SNMP credentials necessary for probing devices using SNMP.</p>	<p>Queries and collects system OIDs such as SysDescr and sysUpTime.</p>
CLI (Device Command-Line by Telnet or SSH)	<ul style="list-style-type: none"> <li>• Similar data set to SNMP</li> <li>• May be used instead of, or in combination with, SNMP</li> </ul>	<p>Requires correctly defined admin login tuples and Enable passwords where needed for device types.</p> <p>You may test credentials against devices and assign CLI credentials to individual objects, overriding Grid-level and Network-level credential settings.</p>	<p>Uses standard device-language scripts and configured Telnet or SSH connection settings to collect discovery data.</p>

Discovery Type	Returned Data	Guideline	Mechanism
vDiscovery	<ul style="list-style-type: none"> <li>• IP addresses</li> <li>• MAC addresses</li> <li>• OS</li> <li>• Discovered name</li> <li>• Virtual entity type</li> <li>• Virtual entity name</li> <li>• Virtual cluster</li> <li>• Virtual datacenter</li> <li>• Virtual switch</li> <li>• Virtual host</li> <li>• Virtual host adapter</li> </ul>	<p>Add the VMware vSphere servers on which you want to perform the vDiscovery.</p> <p>For information about how execute a vDiscovery, see <a href="#">Configuring vDiscovery Jobs</a>.</p>	<p>The appliance communicates with the vSphere servers to collect discovery data on virtual machine instances.</p>

## Starting and Stopping the Discovery Service



### Note

The discovery service can only be started on Grid members that are configured as the Consolidator or Probe (i.e. the Grid members with valid Discovery installed).

Each discovery member requires separate **Discovery** licenses, and must have a running discovery service. Consider the following before starting or stopping the discovery service:

- The Grid Master does not run the discovery service.
- Appliances running a Discovery license and the discovery service do not support HA pairs.
- Discovery Probe appliances appear as Grid members in Grid Manager.
- All appliances running discovery must have the Discovery license installed before starting the service.
- Appliances running discovery do not run core network services such as DNS and DHCP. Discovery appliances may also run the NTP service.
- If you expect to run a single appliance in the Grid for discovery, the appliance is designated as a Consolidator, and also performs Probe discovery operations.
- When you add a new Grid member with a Discovery license, the appliance is set automatically to the following:
  - A Consolidator, if no other discovery member exists in the Grid.
  - A Probe, when at least one discovery appliance exists in the Grid



#### Note

When a member joins the Grid and applies a Discovery license for the first time, the admin user needs to log off and log in again to Grid Manager to see the discovery-enabled functionality.

For information about discovery configuration at the service level, see [Configuring Discovery Properties](#).

## Starting the Discovery Service

To start the discovery service on a licensed Consolidator or Probe appliance:

1. From the **Grid** tab, select the **Grid Manager** tab, and click the **Services** tab.
2. Click the **Discovery** icon to display the list of members running the discovery service.
3. Select the discovery member or members for which you wish to start the service.
4. Expand the Toolbar and click **Start**.  
The appliance asks you to verify that you want to proceed with starting the service for the selected member.
5. Click **Yes**.

## Stopping the Discovery Service

To stop the discovery service on the Consolidator or Probe appliance:

1. From the **Grid** tab, select the **Grid Manager** tab, and click the **Services** tab.
2. Click the **Discovery** icon to display the list of members running the discovery service.
3. Select the discovery member or members for which you wish to stop the service.
4. Expand the Toolbar and click **Stop**.  
The appliance asks you to verify that you want to proceed with stopping the service for the selected member.
5. Click **Yes**.

## Managing Discovered Data

After a discovery, key information is collected and displayed in the following tabs: **Data Management** tab -> **Devices** -> **Interfaces**, **Networks**, **IP Addresses**, and **Assets** tabs of Grid Manager. You can view information about each discovered entity in one of these tabs. For more information, see [Viewing Discovered Devices and their Properties](#).

A discovered entity is considered "unmanaged" if it is discovered in a network for which no information is being stored in the NIOS database. You are not able to configure unmanaged objects in NIOS. Depending on the nature of the discovered entity, you may convert certain unmanaged entities into managed objects so you can manage them through Grid Manager. When an entity is in the managed state, you can configure settings such as applying permissions to it, limiting who can modify the configurations and deployments, and when those changes can be applied. You cannot do so with unmanaged objects.

Grid Manager allows you to convert certain unmanaged devices, interfaces, networks, and assets to the following IPAM object types:

- Host record, as described in [About Host Records](#).
- A and PTR records, as described in [Managing A Records](#) and in [Managing PTR Records](#).
- Fixed IPv4 Address, as described in [Configuring IPv4 Fixed Addresses](#).
- Fixed IPv6 Address, [Configuring IPv6 Fixed Addresses](#).

When converting unmanaged entities to managed objects in NIOS, you can choose to convert them one at a time or as a group. To convert a single entity, just select a specific entity and perform the conversion. To convert multiple entities to the same IPAM object type, you can select the entities you want to manage and then perform a bulk conversion. You can also perform an automatic conversion for unmanaged entities in a network view by configuring conversion rules at the Grid level. For information about how to create conversion rules, see [Configuring Automatic Conversion Rules](#).




#### Note

In a bulk conversion, if one or more of the selected entities are not eligible for conversion because they do not have a discovered name (FQDN) or due to other reasons, Grid Manager displays a warning message indicating only eligible entities are displayed in the conversion wizard and qualified for conversion. Those that cannot be converted are ignored and will remain in unmanaged state.

For more information about how to convert unmanaged entities to managed objects, see the following sections:

- For unmanaged devices, see [Converting Unmanaged Devices to Managed Devices](#) below.
- For unmanaged interfaces, see [Converting Unmanaged Interfaces to Managed Status](#) below.
- For unmanaged networks from the **Data Management -> Devices -> Networks** tab, see [Converting Unmanaged Networks to Managed Status](#) below.
- For unmanaged networks from the **IPAM** or **DHCP** tab, see [Converting Unmanaged Networks under IPAM to Managed Status](#) below.
- For unmanaged IP Addresses, see [Converting Unmanaged IP Addresses to Managed Status](#) below.
- For unmanaged assets, see [Converting Unmanaged Assets to Managed Status](#) below.
- For automatic conversions using global conversion rules, see [About Automatic Conversion Rules](#).

### Converting Unmanaged Devices to Managed Devices



#### Note

For convenience, the home **DataManagement->Devices** page provides a quick filter to list only managed devices. In the Devices page, all devices highlighted in yellow indicate a device that is unmanaged.

Device discovery allows you to define a fully managed state for any discovered routers, switches, firewalls, end hosts, and other network infrastructure devices. The process differs from converting an unmanaged network, because you can bind a discovered device to a fixed address, a PTR Record, a host record or an IPv4 reservation. Doing so offers the following benefits:


- **Host Record** – Infoblox hosts are data objects that contain DNS, DHCP, and IPAM data of the assigned addresses. Host objects allow you to assign multiple IPv4 and IPv6 addresses to a host. When you create a host record, you are specifying the name-to-address and address-to-name mappings for the IP address that you assign to the host. The Infoblox DNS server then uses this data to respond to DNS queries for the host. This establishes the identity of any infrastructure device or asset on the network. For more information, see [About Host Records](#).
- **A Record** – An A (address) record is a DNS resource record that maps a domain name to an IP address. A records essentially tell DNS that a host exists inside a domain, as part of a forward-mapping zone. All traffic to a domain or subdomain is directed to the IP address specified by the A record. The DNS zone must exist in the Grid Manager before attempting to assign A records to devices. For more information about A records, see [Managing A Records](#).
- **PTR Record** – Maps a device IP address to a host name in a reverse-mapping zone. The zone must already be defined before assigning the new PTR record object. For more information about PTR records, see [Managing PTR Records](#).
- **Fixed Address** – A fixed address represents a persistent link between an IP address and one of the following: MAC address, Client identifier, or Circuit ID/remote ID in the DHCP relay agent option (option 82). Most applications in the current context are for a MAC address. You can configure fixed addresses for network devices,

such as routers and printers, that do not frequently move from network to network. By creating fixed addresses for network devices, clients can reliably reach them by their domain names. Some network devices, such as web or FTP servers, can benefit from having fixed addresses for this reason. For more information about these object types, see [Configuring IPv4 Fixed Addresses](#) and [Configuring IPv6 Fixed Addresses](#).

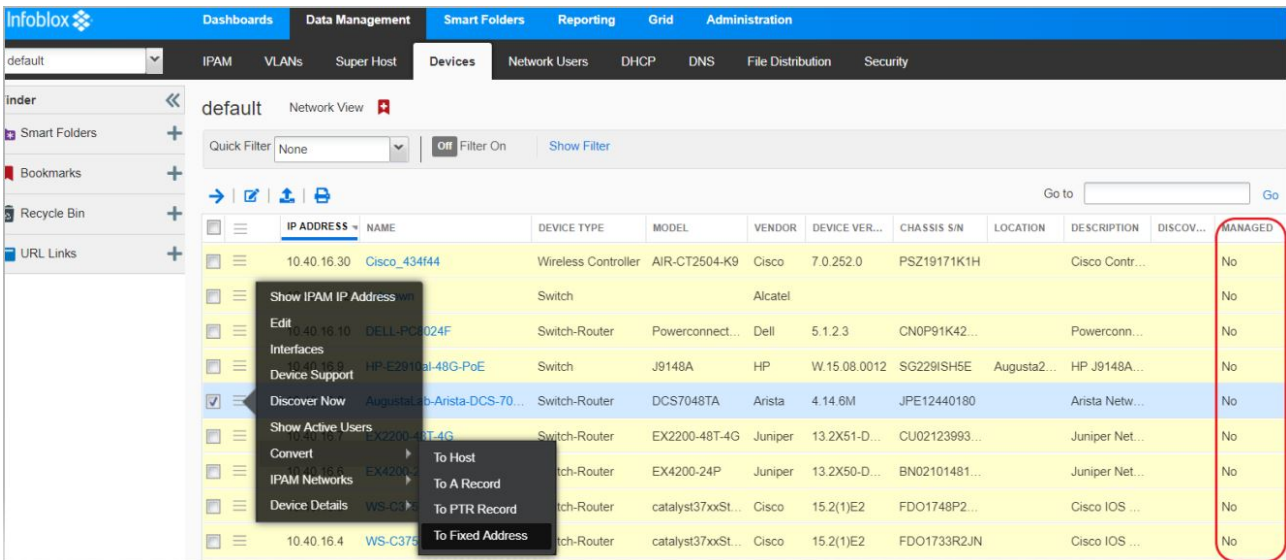
Each object type has its own characteristics that you may apply to specific types of discovered devices. For many infrastructure devices such as routers, or Assets such as servers, a fixed address is a likely choice.

Begin by examining the **Data Management → Devices** page. The **Managed** column, as shown in the following figure, can list one of three possible states for all discovered devices:

- **Blank value**—indicates that the device is not known to IPAM, because insufficient information is available to identify and catalog the device at the present time;
- **No**—Shows that the device in the Devices table is not managed under IPAM or Grid Manager, but enough information is collected and the device can be converted to the Managed state.
- **Yes**—The device in the table is a managed object with an IPAM object type (host, A record, PTR record or fixed address).

An Action icon  (shown as a gear in each row for the table) provides the quickest access to the **Convert** feature for unmanaged devices, as shown in the following figure. You can also click **Convert** in the Toolbar to select the object type.

#### Unmanaged Devices That May be Converted to IPAM objects



To convert unmanaged devices to managed objects:

1. From the **Data Management** tab, select the **Devices** tab.
2. Click the **Next Page** and **Last Page** icons to locate the device you want to convert.
3. To convert a single device: Click the Action icon

next to the device you want to convert (this automatically selects it), and then select **Convert → To Host, To A Record, To PTR Record, or To Fixed Address** from the menu.

To convert multiple devices (bulk conversion): Select the checkboxes of the devices you want to convert to the same IPAM object type. From the Toolbar, select **Convert → To Host, To A & PTR Record, or To Fixed Address** from the menu.

4. For a single device: The respective object editor appears based on the conversion type you have selected. For example, if you select **To Host**, the *Host* editor appears. In the editor, define the required **General** settings for the new object. You can also define other settings you need from any of the tabs in the editor. For details about how to configure these settings, refer to the online Help in Grid Manager or see the appropriate sections in this guide. For bulk conversion: The respective bulk conversion wizard (such as the *Convert Unmanaged IP Addresses to Host Record* wizard) appears based on the conversion type you have selected. Note that the Network Insight creates the names of the new managed objects using the Discovered Name (FQDN) for the entities being converted. In the wizard, Grid Manager displays the IP addresses of the selected

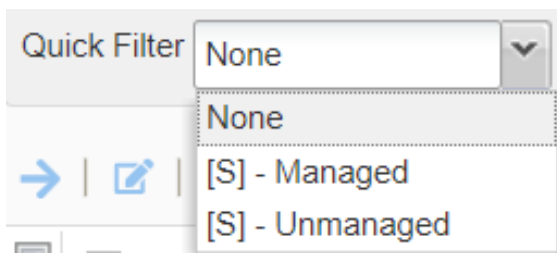
devices that are eligible for conversion in the **Selected IP Addresses** table. Entities that are not eligible for conversion will not be converted and will not appear in this table.

Based on the selected conversion type, complete the following to start a bulk conversion:

- a. **To Host:** You can select **Enable in DNS** and/or **Enable in DHCP** so the appliance can serve DNS and/or DHCP for the selected IP addresses in the host record. When you enable DNS, you must select a DNS zone for all entities that do not have an FQDN.
  - b. **To A & PTR Record:** Network Insight converts the selected entities to A & PTR records simultaneously. You must select a DNS zone for all entities that do not have an FQDN.
  - c. **To Fixed Address:** Network Insight automatically converts all selected IP addresses to fixed addresses in a bulk conversion. As with host record and A/PTR record conversions, entities without a Discovered Name are not eligible for conversions to fixed addresses.
5. In all bulk conversions, you can define extensible attributes for the selected IP addresses. After you configure the necessary settings, you can convert the discovered entities immediately or schedule the conversion by selecting **Later** and entering the date and time.
  6. Click **Save & Close** to make the conversion. You can return to the managed object at any time to make further configuration changes.

In the **Managed** column for all converted entities, Grid Manager displays **Yes** to indicate their managed status. You can now manage these IPAM objects through Grid Manager.

### Converting Unmanaged Interfaces to Managed Status



#### Note

For convenience, a device Interfaces panel provides a quick filter to list only managed interfaces for the device. When a device is converted to managed status, interfaces in the device may remain in unmanaged state. If the interface has an IP address that is recognized under IPAM, it may be converted to managed state.

Interfaces that appear in the Interfaces table for a device may be converted to managed status, under specific circumstances. If an interface is bound to an IP address that is present in an IPAM network (for example, a leaf network inside a network container under IPAM), that interface can be converted to managed status.

For any device, any interface with a hotlink to IPAM may be converted. Examples are shown in the following figure.



## Device Interfaces Available for Conversion

NAME	IP ADDRESS	VRF	NETWORK VIEW	MAC ADDRESS	ACTIVE USERS	PORT TYPE	PORT SPEED	ADMIN STATUS	OPERATION STATUS	TRUNK	DESCRIPTION	STATUS	IPAM TYPE
GigabitEther...		default		E0:89:9D:43:4F:46	0	gigabitEL...	100 Mbps	Up	Up	Off	Unit: 0 Slot: 0 Port: 3...		
GigabitEther...		default		E0:89:9D:43:4F:45	0	gigabitEL...	100 Mbps	Up	Down	Off	Unit: 0 Slot: 0 Port: 2...		
GigabitEther...		default		E0:89:9D:43:4F:47	0	gigabitEL...	100 Mbps	Up	Down	Off	Unit: 0 Slot: 0 Port: 4...		
Virtual Interface	1.1.1.1	default		E0:89:9D:43:4F:44	0	gigabitEL...	1 Gbps	Up	Up	Off	Virtual Interface		
GigabitEther...	10.40.16.30	default		E0:89:9D:43:4F:44	0	gigabitEL...	1 Gbps	Up	Up	Off	Unit: 0 Slot: 0 Port: 1... Used		Unmanaged

In the figure above, two interfaces provide a link to their respective IPAM interface pages, and show an **IPAM Type of Unmanaged**. These ports may be converted to IPAM objects and managed under Grid Manager.

1. From the **Data Management** tab, select the **Devices** tab.
2. Click the **Next Page** and **Last Page** icons to locate the device through which you want to locate the interfaces to convert.
3. Click the **Name** link of the device.
4. Click the **Interfaces** tab for the chosen device. This tab lists all ports discovered on the device.
5. To convert a single interface: Click the Action icon

next to the interface you want to convert (this automatically selects it), and then select **Convert -> To Host, To A Record, To PTR Record, or To Fixed Address** from the menu.

To convert multiple interfaces (bulk conversion): Select the checkboxes of the interfaces you want to convert. From the Toolbar, select **Convert -> To Host, To A & PTR Record, or To Fixed Address** from the menu.

6. For a single interface: The respective object editor appears based on the conversion type you have selected. For example, if you select **To Host**, the *Host* editor appears. In the editor, define the required **General** settings for the new object. You can also define other settings you need from any of the tabs in the editor. For details about how to configure these settings, refer to the online Help in Grid Manager or see the appropriate chapters in this guide. For bulk conversion: The respective bulk conversion wizard (such as the *Convert Unmanaged IP Addresses to Host Record* wizard) appears based on the conversion type you have selected.

Note that the Network Insight creates the names of the new managed objects using the Discovered Name (FQDN) for the entities being converted. In the wizard, Grid Manager displays the IP addresses of the selected devices that are eligible for conversion in the **Selected IP Addresses** table. Entities that are not eligible for conversion will not be converted and will not appear in this table.

Based on the selected conversion type, complete the following to start a bulk conversion:

- a. **To Host:** You can select **Enable in DNS** and/or **Enable in DHCP** so the appliance can serve DNS and/or DHCP for the selected IP addresses in the host record. When you enable DNS for the record, you must select a DNS zone for all entities that do not have an FQDN.
  - b. **To A & PTR Record:** Network Insight converts the selected entities to A & PTR records simultaneously. You must select a DNS zone for all entities that do not have an FQDN.
  - c. **To Fixed Address:** Network Insight automatically converts all selected IP addresses to fixed addresses in a bulk conversion. As with host record and A/PTR record conversions, entities without a Discovered Name FQDN are not eligible for conversions to fixed addresses.
7. In all bulk conversions, you can define extensible attributes for the selected IP addresses. After you have configured the necessary settings, you can convert the discovered entities immediately or schedule the conversion by selecting **Later** and entering the date and time.

7. Click **Save & Close** to make the conversion. The interface is associated with the new IPAM object.

In the **IPAM Type** column for all converted entities, Grid Manager displays **Managed** to indicate their managed status. You can now manage these IPAM objects through Grid Manager.

## Converting Unmanaged Networks to Managed Status

On the **Data Management** → **Devices** page, click a device name, and then open the **Networks** tab. You see the following types of networks based on their managed or unmanaged status:

- **Unmanaged:** These networks are displayed in yellow rows with the value of **No** in the Managed column. Shows that the network is not managed under IPAM/Grid Manager, but enough information is catalogued so that the network can be converted to managed status. You can provision these networks onto devices.
- **Managed:** These networks are displayed in white rows with the value of **Yes** in the Managed column. These networks are currently managed under IPAM, converted to an IPAM network. You can provision and de-provision managed networks.
- **The so called "non-NIOS networks":** These networks are displayed in grey rows with a blank value in the Managed column. Indicates that the network is discovered but available information is insufficient to identify and catalog the network in IPAM at the present time. This can be due to the following reasons:
  - The admin or operation status of the corresponding interface is "down". That is, the interface is either disconnected physically or disabled by administrator.
  - The prefix length for the network is /32 (ipv4) or /128 (ipv6). Network Insight treats this as a route to a specific device rather than a subnet, therefore it does not create such network in IPAM.
  - The route for this interface is configured incorrectly.
  - The route is learned from a remote source via BGP, OSPF and so on (i.e., indirect next hop), or come from a static route using the netmgmt protocol. Network Insight creates networks in IPAM for only direct and local routes from routing tables.
  - The network is within a VRF and the VRF is not mapped to a network view. VRF mapping is required in this case for the network to appear in IPAM. Some time after the VRF is mapped, the network turns from non-NIOS to unmanaged, or managed if the network is already present in IPAM.

Unmanaged networks listed under discovered devices present the same conversion features as networks listed under IPAM (see the following section [Converting Unmanaged Networks under IPAM to Managed Status](#) below).

Converting a network in the device context is the same as converting it at the top IPAM level. You cannot apply services or IPAM objects to IP addresses in unmanaged networks until the networks are converted to managed status. Many operations cannot be carried out on unmanaged networks, including editing, splitting, resizing, permissions changes and many other tasks.

1. From the **Data Management** tab, select the **Devices** tab.
2. Click the **Next Page** and **Last Page** icons to locate the device through which you want to locate the assets to convert.
3. Click the device name.
4. Click the **Networks** tab.
5. To convert a single network: Click the Action icon

next to the network you want to convert to the Managed state (this automatically selects it), and then select **Convert** from the menu.

To convert multiple networks (bulk conversion): Select the checkboxes of the networks you want to convert to the Managed state. From the Toolbar, select **Convert** from the menu.

6. The *Network* editor appears. In the editor, define the required **General** settings for the network. You can also define other settings you need from any of the tabs in the editor. For details about how to configure these settings, refer to the online Help in Grid Manager.  
Note Networks inherit discovery setting from their parent networks. Discovery will be disabled for networks that do not have a parent network.
1. If necessary, select the **Disable for DHCP** checkbox. When you convert a network to managed status, Grid Manager uses the discovered Router IP for the network to automatically populate the Router IP value in DHCP configurations for the selected network. Selecting this option disallows the converted network from being usable under DHCP.
2. If necessary, click the **Discovery** tab and click **Enable Discovery** to start discovery on the network after it is converted to Managed state. You can also elect not to discover the network, by leaving the checkbox clear.
  - Click **Select Member** to choose the Probe member through which the network may be discovered.
  - If necessary, click **Override** under **Polling Options**, and modify the device discovery polling options for the network. You can also specify **Discovery Exclusions** or a **Discovery Blackout** period.

- A number of associated DNS and DHCP services are also available for configuration for the new IPAM network, including DHCP Forwarding, IPv4 DDNS settings, and an array of other DDI service settings for the network. None of these configurations are required in order for the network to be in Managed state, but may be required for other purposes.

7. Click **Save & Close** to make the conversion.

In the **Managed** column in the **Networks** tab, Grid Manager displays **Yes** for all converted entities to indicate their Managed status. You can now manage the networks through Grid Manager.

## Converting Unmanaged Networks under IPAM to Managed Status



### Note

When you convert a network to managed status, Grid Manager uses the discovered Router IP for the network to automatically populate the Router IP value in DHCP configurations for the selected network. Conversion for DHCP is optional; you can choose to **Disable for DHCP** when you convert the network.

The IPAM tab lists all discovered networks as unmanaged, highlighted in yellow. Administrators cannot apply services or IPAM objects to IP addresses in unmanaged networks until the networks are converted to managed status. You can explore unmanaged networks through the IP Map and IP List views, but many operations cannot be carried out on unmanaged networks, including editing, splitting, resizing, permissions changes and other tasks.



### Note

Unmanaged IP addresses that are part of an unmanaged network cannot be independently converted to a managed IP address.

Unmanaged networks can be converted at the IPAM main page and at the device level under **Data Management** → **Devices**, selecting a device and opening the **Networks** page.

Under IPAM, the **Managed** column for the Network tables can show one of two possible states for all discovered IPAM networks:

- **No**—Shows that the network is not managed under IPAM/Grid Manager, but enough information is catalogued that the device can be converted to Managed state. This state is required before a network can be converted to managed status.
- **Yes**—The network shown in the table is now Managed under IPAM, converted to an IPAM network.

You can discover the network again after it is converted, or keep discovery disabled and execute it at another time, or impose blackout periods that limit the time windows under which discovery can execute on the network. Other management benefits include the ability to enable Infoblox services to the network;

1. From the **Data Management** tab, select the **IPAM** tab.
2. To convert a single network: Click the Action icon

next to the network you want to convert to the Managed state (this automatically selects it), and then select **Convert** from the menu.

To convert multiple networks (bulk conversion): Select the checkboxes of the networks you want to convert to the Managed state. From the Toolbar, select **Convert** from the menu.

You do not need to take further action other than **Save & Close** to set the network as Managed.

3. If necessary, you can select the converted network and do the following in the *Network* editor:
4. Select the **Disable for DHCP** checkbox to disallow the converted network from being usable under DHCP.
  - If necessary, click the **Discovery** tab and click **Enable Discovery** to start discovery on the network immediately after it is converted to Managed state. You can also elect not to discover the network, by leaving the checkbox clear.
  - Click **Select Member** to choose the Probe member through which the network may be discovered.
  - If necessary, click **Override** under **Polling Options**, and modify the device discovery polling options for the network. You can also specify **Discovery Exclusions**, port control settings, or a **Discovery Blackout** period.

- A number of associated DNS and DHCP services are also available for configuration for the new IPAM network, including DHCP Forwarding, IPv4 DDNS settings, and an array of other DDI service settings for the network. None of these configurations are required in order for the network to be in Managed state, but may be required for other purposes.
5. The IPAM main page shows **Yes** as the value for the network under the **Managed** column. The network is now under management of IPAM and can participate in Infoblox services.



#### Note

Most conversion operations for networks and individual IP addresses are managed under IPAM and are described in the section [Managing IPv4 and IPv6 Addresses](#).

## Converting Unmanaged IP Addresses to Managed Status

The principle for converting an IP address to managed status is the same for IP addresses as for interfaces. If an IP address is unmanaged but is present in IPAM, the IP can be converted to managed status.

Figure 15.6 Converting an IP Address to Managed Status

The screenshot shows the IPAM interface with the 'Data Management' tab selected. The 'Devices' sub-tab is active, showing a list of IP addresses for a specific device. A context menu is open over the IP address E0:89:9D:43:4F:44, with the 'Convert -> To Host' option highlighted. The table below shows the details of the IP addresses.

NAME	IP ADDRESS	VRF...	NETWORK VIEW	V...	MAC ADDRESS	ACTIVE USERS	V...	V...	PORT TYPE	PORT SPEED	ADMIN STATUS	OPERATION STATUS	TRUNK...	DESCRIPTION	STATUS	IPAM TYPE
GigabitEther...			default		E0:89:9D:43:4F:46	0			gigabitET...	100 Mbps	Up	Up	Off	Unit: 0 Slot: 0 Port: 3...		
GigabitEther...			default		E0:89:9D:43:4F:45	0			gigabitET...	100 Mbps	Up	Down	Off	Unit: 0 Slot: 0 Port: 2...		
GigabitEther...			default		E0:89:9D:43:4F:47	0			gigabitET...	100 Mbps	Up	Down	Off	Unit: 0 Slot: 0 Port: 4...		
GigabitEther...	1.1.1.1		default		E0:89:9D:43:4F:44	0			gigabitET...	1 Gbps	Up	Up	Off	Virtual Interface		
			default		E0:89:9D:43:4F:44	0			gigabitET...	1 Gbps	Up	Up	Off	Unit: 0 Slot: 0 Port: 1...	Used	Unmanage

1. From the **Data Management** tab, select the **Devices** tab.
2. Click the **Next Page** and **Last Page** icons to locate the device through which you want to locate the IP address to convert.
3. Click the Name of the device.
4. Click the **IP Addresses** tab for the chosen device. The IP address page lists all ports discovered on the selected device.
5. To convert a single IP address: Click the Action icon next to the IP address you want to convert (this automatically selects it), and then select **Convert -> To Host**, **To A Record**, **To PTR Record**, or **To Fixed Address** from the menu.  
 To convert multiple IP addresses (bulk conversion): Select the checkboxes of the IP addresses you want to convert. From the Toolbar, select **Convert -> To Host**, **To A & PTR Record**, or **To Fixed Address** from the menu.
6. For a single IP address: The respective object editor appears based on the conversion type you have selected. For example, if you select **ToHost**, the *Host* editor appears. In the editor, define the required **General** settings for the new object. You can also define other settings you need from any of the tabs in the editor. For details about how to configure these settings, refer to the online Help in Grid Manager or see the appropriate chapters in this guide.  
For bulk conversion: The respective bulk conversion wizard (such as the

*Convert Unmanaged IP Addresses to Host Record* wizard) appears based on the conversion type you have selected.

Note Network Insight creates the names of the new managed objects using the Discovered Name (FQDN) for the entities being converted. In the wizard, Grid Manager displays the IP addresses of the selected devices that are eligible for conversion in the **Selected IP Addresses** table. Entities that are not eligible for conversion will not be converted and will not appear in this table.

Based on the selected conversion type, complete the following to start a bulk conversion:

- a. **To Host:** You can select **Enable in DNS** and/or **Enable in DHCP** so the appliance can serve DNS and/or DHCP for the selected IP addresses in the host record. When you enable DNS for the record, you must select a DNS zone for all entities that do not have an FQDN.
- b. **To A&PTR Record:** Network Insight converts the selected entities to A & PTR records simultaneously. You must select a DNS zone for all entities that do not have an FQDN.
- c. **To Fixed Address:** Network Insight automatically converts all selected IP addresses to fixed addresses in a bulk conversion. As with host record and A/PTR record conversions, entities without a Discovered Name FQDN are not eligible for conversions to fixed addresses.

In all bulk conversions, you can define extensible attributes for the selected IP addresses. After you have configured the necessary settings, you can convert the discovered entities immediately or schedule the conversion by selecting **Later** and entering the date and time.

Click **Save&Close** to make the conversion. You can return to the managed object at any time to make further configuration changes.

7. In the **Managed** column for all converted entities, Grid Manager displays **Yes** to indicate the Managed status. The selected IP addresses are associated with the newly converted IPAM objects. You can now manage these IPAM objects through Grid Manager.

## Converting Unmanaged Assets to Managed Status

1. From the **Data Management** tab, select the **Devices** tab.
2. Click the **Next Page** and **Last Page** icons to locate the device through which you want to locate the assets to convert.
3. Click the Name of the device.
4. Click the **Assets** tab for the chosen device. The Assets page lists all network devices that have been detected through connections to the current device. If the device is a switch, you may see a list of end hosts and network servers in the list, along with any other network infrastructure devices that are neighboring the current device.
5. To convert a single asset: Click the Action icon

next to the asset you want to convert (this automatically selects it), and then select **Convert → To Host, To A Record, To PTR Record, or To Fixed Address** from the menu.

To convert multiple assets (bulk conversion): Select the checkboxes of the assets you want to convert. From the Toolbar, select **Convert -> To Host, To A & PTR Record, or To Fixed Address** from the menu.

6. For a single asset: The respective object editor appears based on the conversion type you have selected. For example, if you select **To Host**, the *Host* editor appears. In the editor, define the required **General** settings for the new object. You can also define other settings you need from any of the tabs in the editor. For details about how to configure these settings, refer to the online Help in Grid Manager or see the appropriate chapters in this guide.

For bulk conversion: The respective bulk conversion wizard (such as the *Convert Unmanaged IP Addresses to Host Record* wizard) appears based on the conversion type you have selected.

Note Network Insight creates the names of the new managed objects using the Discovered Name (FQDN) for the entities being converted. In the wizard, Grid Manager displays the IP addresses of the selected devices that are eligible for conversion in the **Selected IP Addresses** table. Entities that are not eligible for conversion will not be converted and will not appear in this table.

Based on the selected conversion type, complete the following to start a bulk conversion:

- a. **To Host:** You can select **Enable in DNS** and/or **Enable in DHCP** so the appliance can serve DNS and/or DHCP for the selected IP addresses in the host record. When you enable DNS for the record, you must select a DNS zone for all entities that do not have an FQDN.
- b. **To A & PTR Record:** Network Insight converts the selected entities to A & PTR records simultaneously. You must select a DNS zone for all entities that do not have an FQDN.
- c. **To Fixed Address:** Network Insight automatically converts all selected assets to fixed addresses in a bulk conversion. As with host record and A/PTR record conversions, entities without a Discovered Name FQDN are not eligible for conversions to fixed addresses.

In all bulk conversions, you can define extensible attributes for the selected assets. After you have configured the necessary settings, you can convert the discovered entities immediately or schedule the conversion for a later date by selecting **Later** and entering the date and time.

7. Click **Save & Close** to make the conversion. You can return to the managed object at any time to make further configuration changes.

In the **Managed** column for all converted entities, Grid Manager displays **Yes** to indicate the managed status. The selected assets are associated with the newly converted IPAM objects. You can now manage these IPAM objects through Grid Manager.

## Configuring Automatic VRF Mapping

You can configure Network Insight to automatically assign network views to VRFs by defining mapping rules that match unassigned VRFs—VRFs that do not have associated network views. You can also disable the automatic VRF mapping and manually assign network views to discovered VRFs, as described in [Viewing Discovered VRFs and Mapping Network Views](#).

When you enable automatic VRF mapping, you can add VRF mapping rules that Network Insight uses to map network views to VRFs that do not already have an assigned network view. When you configure mapping rules, you can define criteria using regular expressions and place the rules in specific order to define their priorities. Network Insight matches the rule criteria to the discovered VRF names, starting with the first rule in the VRF Mapping Rules table. You also have a choice to map an unassigned VRF to the network view from which one of the interfaces the VRF is reached, if none of the mapping rules match the VRF name.

To configure automatic VRF mapping and mapping rules for unassigned VRFs, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click the **Discovery** service.
2. Click **Edit** -> **Grid Discovery Properties** from the Toolbar.
3. In the *Grid Discovery Properties* editor, select the **VRF Mapping Rules** tab, and then complete the following:
  - **Enable the automatic VRF mapping rules defined below for unassigned VRFs:** Select this to enable automatic VRF mapping so you can define mapping rules that Network Insight uses to map network views to unassigned VRFs that match the criteria of the rules.
  - **Enable the automatic VRF mapping rules and system mapping extensions:** Select this to enable the VRF Mapping Rules table so you can define mapping rules that Network Insight uses to map network views to unassigned VRFs that match the criteria of the rules; and in cases where none of the rules match a VRF name, Network Insight maps the VRF to the network view from which one of the interfaces the unassigned VRF is reached.
  - **Disable automatic VRF mapping and only use manually defined VRF mapping:** Select this to disable the VRF Mapping Rules table. When you select this, Network Insight does not perform any evaluation of the VRF mapping rules. You can manually assign or unassign network views to the discovered VRFs, as described in [Viewing Discovered VRFs and Mapping Network Views](#).

When you enable automatic VRF mapping, you can add mapping rules to the VRF Mapping Rules table, as follows:

1. Click the Add icon, and the appliance adds a row to the table.
2. In the table, click each of the following fields and enter the values accordingly:
  - **Network View:** The network view that you want to use for all matching VRFs. You can click this field and select a network view from the drop-down list that displays all the configured network views, including the default network view.
  - **Order:** The order and priority in which Network Insight evaluates the mapping rules. Each time you add a new rule, the appliance automatically appends the rule to the end of the list and assigns the next incremental number to the rule. To reorder the list, you can select a rule and use the up and down arrows next to the table to move the rules to its desired position so you can set the priority for the rule evaluation. Network Insight evaluates the rules based on the order, starting with 1 as the highest priority.
  - **Criteria:** The criteria that Network Insight uses to match the VRF name of an unassigned VRF. You can use POSIX regular expressions to define the mapping criteria. The appliance validates the rule when you save the configuration, and it returns an error message if the criteria is invalid.
  - **Comment:** Enter a comment about the VRF mapping rule. Click the Add icon again to define another mapping rule.
3. Save the configuration.





#### Note

All the VRF mapping rules that are currently configured for the Grid are displayed in the VRF Mapping Rules table.

You can also do the following in the **VRF Mapping Rules** tab:

- Select a specific rule and click the Delete icon to remove it from the table.
- Use the up and down arrows next to the rules table to reorder the rules.
- Use the **Go to** function to search for a specific mapping rule. With the autocomplete feature, you can just enter the first few characters of a network view in the **Go to** field and select the network view from the possible matches.

## About Network Insight

Infoblox Network Insight provides discovery features for detecting and managing devices in your network infrastructure. You can use discovery to collect device data and manage it through Grid Manager. For more information about device management, see [Managing Discovered Data](#).

You can view the operating state of all discovered network infrastructure devices and newly discovered IP networks, including but not limited to routers, wireless routers and access points, firewalls, load balancers, Ethernet L2/L3 switches, end hosts and other devices in end host networks, end host networks, VRF (Virtual Routing and Forwarding) virtual networks, single and multipoint VPNs, SDN and SD-WAN devices, and much more. Network Insight makes it easy to manage and secure your enterprise network by detecting all interfaces for every discovered device and providing specific information about them. For information about discovering VRF virtual networks, see [Discovering VRF Virtual Networks](#).

The Infoblox IPAM feature set also provides control mechanisms by including and excluding networks and IP addresses for discovery. You can schedule and define when discovery takes place on any network, and define blackout periods during which no discovery tasks occur. Infoblox IPAM and DHCP functions also extend network control to assigning of discovered switch and router interfaces to IPAM objects such as IP networks, IP reservations, and host records. The assignments are called *port reservations* and are part of a feature set called *port control*, managed through Grid Manager. For more information, see [Port Control Features in Network Insight](#).

Infoblox provides a few reports in which you can view trending information about the device groups, types of devices, and device IP addresses for the devices that are discovered by Network Insight. For information about these reports, see [Status Dashboard](#).

Deployment of Network Insight discovery requires a separate Discovery license and one or more NIOS appliances dedicated to discovery tasks.

Network Insight performs discovery in the following ways:

- Through specification of seed routers, which inform discovery of the various networks that should be examined and catalogued.
- Through discovery of the various Object types you can create under IPAM and DHCP, including IPv4 and IPv6 Networks, Host Records, IPv4 reservations, DHCP ranges, and IPv4/IPv6 fixed addresses.
- An enhanced VM discovery allowing both scheduled VM discovery and immediate discovery of VMs.

## Supported Appliances for Network Insight

Network Insight is supported on the following physical appliances: ND-805, ND-1405, ND-2205, and ND-4005; and the following virtual appliances: ND-805, ND-V1405, and ND-V2205. All appliances that perform discovery require a Discovery license. Appliances with this license only perform discovery tasks and do not perform core DNS or DHCP functions.

## Configuring Super Hosts

A super host is a collection of resource records or fixed addresses that are related. It can contain resource records or fixed addresses that belong to a single network device, such as a router or a switch, or an application server. With super host, you can configure and manage multiple interfaces, IP addresses, and DNS and DHCP records that are associated with the same physical or virtual device. A super host gives you the flexibility to group related objects into a single entity and manage them at one place.

A super host can contain the following record types:

- DNS records - A, AAAA, PTR, and host addresses
- DHCP records - IPv4 or IPv6 fixed addresses.

Note that a super host object and a network view or a DNS view are independent of each other. A super host is an aggregation of NIOS objects that can belong to multiple networks or DNS views, but are related to a single network device. The audit log records all changes to a super host object including addition, modification, deletion, and association or disassociation of resource records.

You can group super host objects in smart folders according to their attributes. For example, you can create a smart folder that contains all records in a specific super host. For more information, see [Smart Folders](#). Use Capacity Report to view the total capacity utilization of super host objects. For more information, see [Using the Capacity Report](#).



### Note

Users must have valid permissions to delete or disable a super host. A resource record can only belong to one super host and cannot be a part of multiple super hosts. You can associate any number of records with a super host.

## Limitations of Super Host

- You can associate extensible attributes with a super host, but it does not support inheritance.
- You cannot associate DNS records that are auto-created with a super host.
- Infoblox has limited control over the resource records, from a zone or a network, which is associated with a super host. This can cause inconsistent behaviors when you enable or disable the resource record from the parent zone or network.

## Administrative Permissions for Super Hosts

Superusers and limited-access users with read/write permission to **All Super Hosts** and corresponding objects can manage super hosts and their associated resource records. By default, the Super Host Admin role supports read/write permission on super hosts.

To define permissions for super hosts:

1. **For an admin group:** From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin\_group* in the Groups table, and then click the Add icon -> **Global Permissions** from the Create New Permission area or select Add -> **Global Permissions** from the Toolbar. For more information, see [About Administrative Permissions](#).
2. **For an admin role:** From the **Administration** tab, select the **Administrators** tab -> **Permissions** tab -> *admin\_role* in the Roles table, and then click Add icon -> **Global Permissions** from the Create New Permission area or select Add -> **Global Permissions** from the Toolbar. For more information, see [About Admin Roles](#).
3. Complete the following in the *Manage Global Permissions* dialog box:
  - **Permission Type:** Select **Super Host Permissions** from the drop-down list.
  - In the table, select **Read/Write**, **Read-Only**, or **Deny** for All Super Hosts.



4. Click **Save & Close**.

You can also define object level permissions for a super host. For more information, see [Defining Object Permissions](#).

## Adding Super Hosts

To add a super host, complete the following:

1. From the **Data Management** tab, select the **Super Host** tab, and then click the Add icon.
2. In the *Add Super Host* wizard, enter the following details:
  - **Name:** Enter a name for the super host.
  - **Comment:** Enter additional information about the super host.
  - **Disabled:** Select this checkbox if you want to disable the super host at this time. Clear the checkbox to enable it. When you disable a super host, the DNS and DHCP records that are associated with it are disabled, but still associated with the super host.
3. Click **Next** to associate DNS or DHCP records with the super host. Click the arrow next to the Add icon and select Add DNS and Host Records or Add DHCP Records. Note that you must have valid DNS and DHCP permissions to add resource records and fixed addresses.
4. The records that you associate with the super host are listed. To delete a specific record, select the respective checkbox and click the Delete icon.
5. Click **Next** to define extensible attributes. For information, see [Managing Extensible Attributes](#).
6. Click **Save & Close**.

## Adding DNS and Host Records

In the *Select Record* dialog box, complete the following to select and associate DNS or host records:

- **Network View:** Select a network view from the drop-down list. This is displayed only if you have configured multiple network views.
- **DNS View:** Select a DNS view from the drop-down list.
- **Zone:** Select a zone from the drop-down list. You must select a zone before you start searching for records. Click **Apply** to apply your settings or **Reset** to reset the settings. Optionally, click the + icon to add another filter rule. To delete an additional filter rule, click the - icon.
- **Find:** Specify the name of the resource record and click **Go** to search for a specific record.
- **Host Record:** Select the checkbox to search for a host record.
- **A Record:** Select the checkbox to search for an A record.
- **AAAA Record:** Select the checkbox to search for an AAAA record.
- **PTR Record:** Select the checkbox to search for a PTR record.
- Click the right arrow to navigate to the search results in the next page.

The appliance displays the following details:

- **Name:** Displays the name of the resource record.
- **IP address:** Displays the IP address of the resource record.
- **Type:** Displays the type of the resource record.
- **Comment:** Displays additional comments about the resource record.
- **Site:** Displays the values that were entered for this pre-defined attribute.
- Click on a row and click **OK** to associate the resource record with the super host, or click **Close** to cancel. To select and associate multiple records, use Shift+click and Ctrl+click.

You can edit a resource record that is associated with a super host. For more information, see [Copying and Modifying Host Records](#).

## Adding DHCP Records

In the *DHCP Objects Selector* dialog box, complete the following to select and associate a DHCP record:

- **Network View:** Select a network view from the drop-down list. This is displayed only if you have configured multiple network views.
- **Network:** Select a network from the drop-down list. You must select a zone before you start searching for records. Click **Apply** to apply your settings or **Reset** to reset the settings. Optionally, click the + icon to add another filter rule. To delete an additional filter rule, click the - icon.
- **Find:** Specify the fixed address and click **Go** to search for a specific record.
- Click the right arrow to navigate to the search results in the next page.

The appliance displays the following details:

- **IP address:** Displays the IP address of the resource record.
- **Name:** Displays the respective IPv4/IPv6 fixed address or host address.
- **Type:** Displays the type of the resource record.
- **Comment:** Displays additional comments about the fixed address.
- Click on the row and click **OK** to associate the resource record with the super host, or click **Close** to cancel. To select and associate multiple records, use Shift+click and Ctrl+click.

When you search for DHCP records, a host record with multiple host addresses will display all the associated host addresses in the *DHCP Objects Selector* dialog box. When you select multiple hosts, the NIOS appliance considers the records as duplicate. Infoblox recommends that you delete the duplicate records before you save them.


You can edit an IPv4/IPv6 fixed address that is associated with a super host. For more information about modifying IPv4 and IPv6 fixed addresses, see [Configuring IPv4 Networks](#) and [Configuring IPv6 Fixed Addresses](#) respectively.

## Viewing Super Hosts

The super host Home panel lists all the super hosts that are configured in the Grid. You can provision new super hosts objects and manage them. To view the configured super hosts, navigate to the **Data Management** tab → **Super Host** tab. Grid Manager displays the following information in the super hosts Home panel:

- **Name:** The name of the super host.
- **Comment:** The information you entered about the super host.
- **Disabled:** Indicates whether the super host is disabled.

You can also do the following in this panel:

- Using the Action icon  , you can modify or delete a super host.
- View associated records. For more information, see [Viewing Resource Records Associated with a Super Host](#) below.
- Sort the data in **ascending** or **descending** order by column. For information about customizing tables in Grid Manager, see [Customizing Tables](#).
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches. For more information, see [Using the Go To Function](#).
- Create a quick filter to save frequently used filter criteria. For more information, see [Using Quick Filters](#).
- Use **Global Search** to search for a specific super host. For information, see [Using Global Search](#).

## Viewing Resource Records Associated with a Super Host

You can view the resource records that are associated with a super host by navigating to the **Data Management** tab → **Super Host** tab, click the respective super host. For each super host object, the panel displays the following by default:

- **Name:** The name of the associated DNS or DHCP resource record.
- **Parent:** The name of the super host object in which the record resides.
- **Type:** The record type. When you search for an unspecified record type, all the associated records are displayed.
- **Data:** Value of a resource record. For example, IP address of an A record.
- **Comment:** Comments that were entered for the resource record.
- **Disabled:** Indicates if the associated DNS/DHCP resource record is disabled.

- **Network View:** The network view associated with the resource record.
- **DNS View:** The DNS view to which the zones belong.
- **Creation Timestamp:** The time at which the DNS resource record was created.

You can also do the following in this panel:

- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches. For more information, see [Using the Go To Function](#).
- Create a quick filter to save frequently used filter criteria. For more information, see [Using Quick Filters](#).
- Click the right arrow to return to the parent object:
  - DNS records, such as A/AAAA/PTR records, and host addresses return to the respective parent zone.
  - DHCP records, such as IPv4 and IPv6 fixed addresses, return to the respective parent network.
- Click the Print icon to print the list of associated resource records.

## Modifying Super Hosts

You can select supported DNS or DHCP records and associate or dissociate them from a super host. Infoblox allows you to update the name, comment, and enable or disable the respective super host. To modify a super host:

1. From the **Data Management** tab, select the **Super Host** tab.
2. Select a super host that you want to modify and click the Edit icon, or click the Action icon next to the respective super host and select **Edit** from the menu.
3. The *Super Host* editor provides the following tabs from which you can modify data:
  - **General** tab: Modify certain general properties. For more information, see [Adding Super Hosts](#) above.
  - **Extensible Attributes** tab: Add or modify values of extensible attributes. For information, see [Managing Extensible Attributes](#).
  - **Permissions** tab: Modify the administrative permissions. For more information, see [About Administrative Permissions for Super Hosts](#) above.
4. Click **Save & Close**.

## Deleting Super Hosts

When you delete a super host, you have an option to delete only the respective super host or delete the associated records also. The deleted super hosts and associated records are moved to the Recycle Bin, from which you can restore or permanently delete them. For information about the Recycle Bin, see [Using the Recycle Bin](#).

To delete a super host:

1. From the **Data Management** tab, select the **Super Host** tab.
2. Select a super host that you want to delete and click the Delete icon, or click the Action icon next to the respective super host and select **Delete** from the menu.
3. The appliance displays the *Delete Confirmation* dialog box to confirm that you want to delete a super host. You can choose to delete the records that are associated with the super host by selecting the **Delete associated records with the Superhost** checkbox. You cannot delete associated records if you do not have write permissions on those objects. When you restore this super host from the Recycle Bin, the associated resource records are also restored.  
If you do not select the **Delete associated records with the Superhost** checkbox while deleting a super host, the associated records will not be deleted. When you restore the deleted super host, only the respective super host is restored and you must associate the resource records manually.

## DNS

This section describes how to configure the Grid to provide DNS services. It includes the following topics:

- [Infoblox DNS Service](#)

- [Configuring DNS Services](#)
- [DNS Views](#)
- [Configuring DNS Zones](#)
- [Configuring DNS Resource Records](#)
- [Configuring DDNS Updates](#)
- [Configuring DNSSEC](#)
- [Managing DNS Traffic Control](#)
- [Configuring IP Routing Options](#)

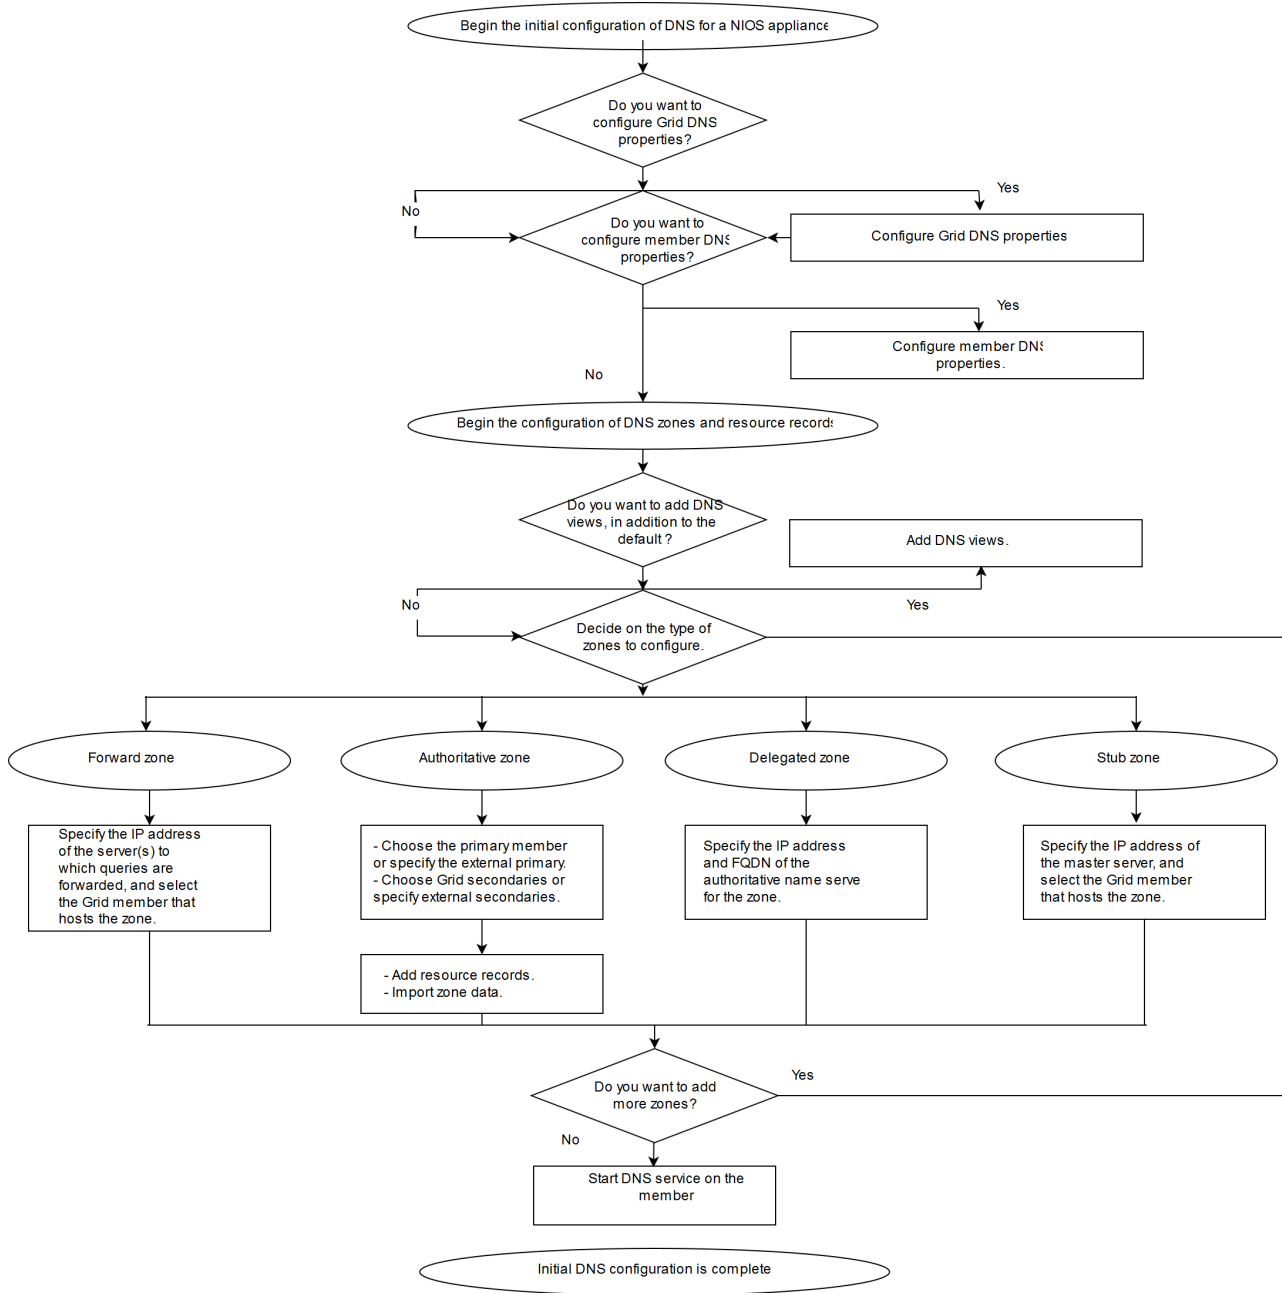
## Infoblox DNS Service

The NIOS appliance uses a standard, BIND-based DNS protocol engine. It interoperates with any other name server that complies with the DNS RFCs. This section provides an overview of the DNS configuration tasks. It includes the following topics:

- [Configuring DNS Overview](#)
- [Inheriting DNS Properties](#)
- [Understanding DNS for IPv6](#)

## Configuring DNS Overview

An overview of the DNS configuration process is outlined in the following diagram, illustrating the required steps for preparing a NIOS appliance for use:



## DNS Configuration Checklist

The following checklist includes the major steps for configuring DNS:

### DNS Configuration Checklist

Step	For more information
Decide if you want to configure DNS properties for the Grid and for individual members	<a href="#">Infoblox DNS Service</a>
Decide if you want to create a new DNS view, in addition to the default DNS view	<a href="#">DNS Views</a>

Step	For more information
Decide which type of DNS zone you want to configure	<a href="#">Configuring DNS Zones</a>
Add hosts and resource records	<a href="#">Configuring DNS Cache Acceleration</a>
Import zone data	<a href="#">Importing Zone Data</a>
Enable DNS service on the member	<a href="#">Starting and Stopping the Discovery Service</a>

## Inheriting DNS Properties

You can configure DNS properties at the Grid, member, zone, and resource records level. The NIOS appliance applies the properties hierarchically, with the Grid at the top of the hierarchy. Grid settings apply to all members in the Grid, unless you override them at the member, zone, or resource record level. When you set DNS properties for a particular member, these properties override the Grid properties and apply to all zones served by that member. When you set properties for a specific zone, they override the member properties and apply to the resource records in the zone. You can also override the zone properties and set properties for specific resource records.

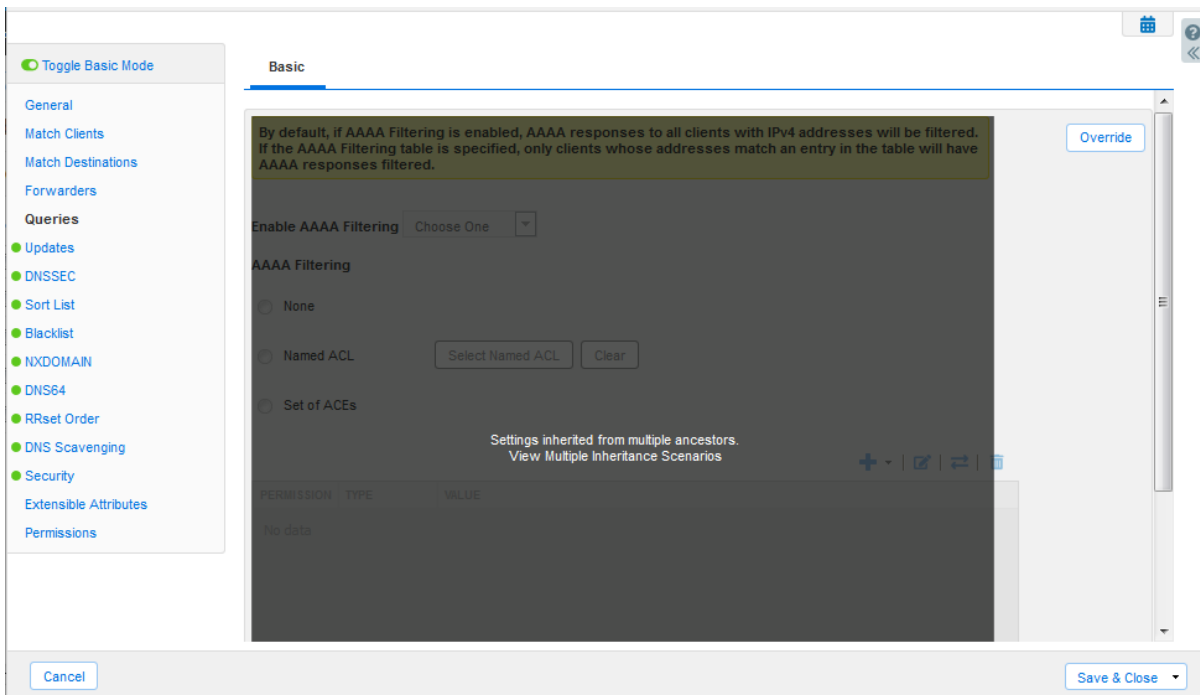
When you configure DNS properties that contain inherited values, the appliance displays the information based on the inheritance sources. There may be times when an object can inherit properties from different sources with different settings. The following table summarizes what the appliance can display:

When you see...	it means...
<b>Inherited From &lt;object&gt;</b>	the DNS property has a definite value from an inheritance source.
<b>Inherited From Upper Level</b>	the appliance cannot yet determine the inherited value or inheritance source for the DNS property.
<b>Inherited From Multiple</b>	the DNS property has the same value that it inherits from multiple sources.
<b>Settings Inherited from Multiple Ancestors, View Multiple Inheritance Scenarios</b>	the DNS property has different values that it inherits from multiple sources, and you can view the values and their corresponding sources by clicking the <b>View Multiple Inheritance Scenarios</b> link.

Based on the information provided, you can then decide whether to override or keep the inherited values. You must have read/write permissions to the DNS resources to override inherited values. You can only view inherited values and paths if you have at least read-only permissions.

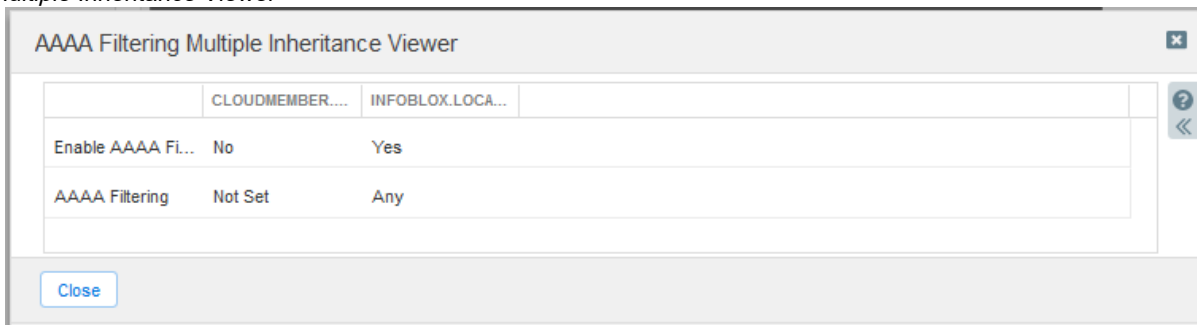
In the example in the below figure, the DNS zone is served by members with different query settings.

*DNS Zone with Different Inherited Settings*



The Multiple Inheritance Viewer indicates that the two servers have different query ACLs, as shown in Multiple Inheritance Viewer figure below. You can then view the Query properties of each member and edit them, or override the setting and specify values that apply to the zone only.

#### Multiple Inheritance Viewer



### Overriding DNS Properties

DNS properties configured at the Grid level apply to the entire Grid. You can choose to keep the inherited properties or override them when you configure the properties for a member, zone, or resource record.

To override an inherited value:

1. In a wizard or editor, click **Override** next to a property to enable the configuration. The **Override** button changes to **Inherit**.
2. Enter a new value to override the inherited value.

### Understanding DNS for IPv6

You can configure NIOS appliances to provide DNS services over IPv4 (Internet Protocol version 4) and IPv6 (Internet Protocol version 6) networks. You can configure the Grid member as IPv6, or dual-mode name server. Dual-mode name server is capable of sending and receiving IPv4 and IPv6 queries and responses. It can serve DNS data in response to both IPv4 and IPv6 queries. The appliance supports authoritative forward-mapping zones containing AAAA records mapping host names to IPv6 addresses, as well as authoritative reverse-mapping zones with PTR records mapping IPv6 addresses to host names.

Configuring a Grid containing an IPv4 primary server and IPv6 secondary servers is not supported. You must enable

IPv6 on both the primary and secondary servers within the Grid to enable them to communicate with each other. Infoblox highly recommends that you enable IPv6 on your Grid appliances before configuring IPv6 authoritative zones. The NIOS appliance supports IPv6 configuration on multiple interfaces, such as LAN1, LAN2, MGMT, anycast (OSPF and BGP), and loopback interfaces. Infoblox integrates IPv6 address management into many of the same places where IPv4 addresses are entered. Data validation occurs on all IP address fields and automatic validation is done to ensure proper entry of either an IPv4 address or an IPv6 address.

The NIOS appliance supports the following DNS functions for IPv6:

- AAAA records—You can import, serve queries, display, add, delete, and modify AAAA records on the appliance. An AAAA record is equivalent to an IPv4 A record, relying upon a forward-mapping zone to map a hostname to an IPv6 address. A single forward-mapping zone can map names to both IPv4 and IPv6 addresses. The appliance autogenerates AAAA records for any of its interfaces that have IPv6 addresses.
- Hosts—You can configure IPv4 and IPv6 addresses for hosts. For information, see [Adding Host Records](#).
- ip6.arpa— A specific domain for IPv6 is used for DNS reverse lookups called ip6.arpa. This domain maps an IPv6 address to a hostname. When you specify an IPv6 network, the appliance automatically creates the appropriate zone under ip6.arpa.
- PTR records—Import, serve queries, display, add, delete, and modify PTR records within an ip6.arpa reverse zone. The PTR record returns a domain name corresponding to an IPv6 address contained in the ip6.arpa zone. The appliance does not autogenerate PTR records; the user must configure PTR records manually.
- DDNS—The appliance supports AAAA and PTR records for DDNS (Dynamic DNS).

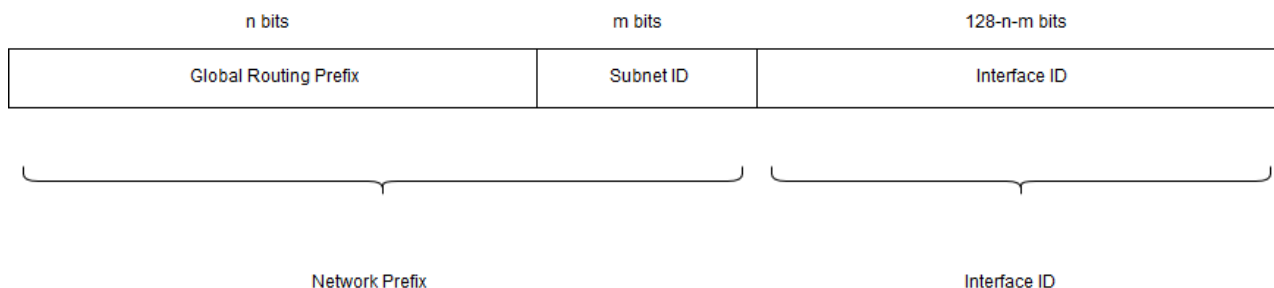
For more information about DNS for IPv6, see RFC 3596, *DNS Extensions to Support IP Version 6*.

## Address Structures

IPv4 uses a 32-bit, 4-octet (each octet separated by decimals) addressing structure to designate sources and destinations within a network. Since there are 32 bits that make up the address, IPv4 can support up to 4 billion unique addresses.

An IPv6 address is a 128-bit number in colon hexadecimal notation. It consists of eight groups of four hexadecimal digits separated by colons (example: 12ab:0000:0000:0123:4567:89ab:0000:cdef). Since there are 128 bits that make up the address, IPv6 can support up to 3.4x10<sup>38</sup> unique addresses. The increase in the number of unique IPv6 addresses is one of the biggest advantages of an IPv6 implementation.

### IPv6 Address Structure



The IPv6 address structure consists of the following:

- Global Routing Prefix—Global routing prefix is a (typically hierarchically-structured) value assigned to a site.
- Subnet ID—Subnet ID is an identifier of a link within the site.
- Interface ID—Interface Identifier. This portion of the address identifies the interface on the subnet. This is equivalent to the host identifier for IPv4 addresses.

When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2006:0000:0000:0123:4567:89ab:0000:cdef can be shortened to 2006::123:4567:89ab:0:cdef. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The NIOS appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered.



## Configuring IPv6 on a Grid Member

You can configure a Grid member in either IPv6 or dual mode (IPv4 and IPv6) to support IPv6 connections. A dual mode Grid member supports both IPv4 and IPv6 connections by configuring an IPv6 address on the member, in addition to the standard IPv4 address.

When you enable IPv6 on a member, you can manually enter the IPv6 gateway address or enable the member to automatically acquire the address from router advertisements. Routers periodically send router advertisements that contain link-layer addresses and configuration parameters. A NIOS appliance that supports IPv6 can listen for router advertisements and obtain the default gateway IP address and link MTU (maximum transmission unit). The link MTU is the maximum packet size, in octets, that can be conveyed in one transmission unit over a link. Thus you can set parameters on a router once and automatically propagate it to all attached hosts.

To configure the member to support IPv6:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.
2. Expand the Toolbar and click **Add** -> **Add Grid Member**.
3. In the *Add Grid Member* wizard, enter the following and click **Next**:
  - **Member Type**: Specify the appliance type of the Grid member. If the member is an Infoblox appliance, select **Infoblox**, which is the default. If the member is a NIOS virtual appliance, select **Virtual NIOS**.
  - **Host Name**: Type the FQDN (fully qualified domain name) of the appliance that you are adding to the Grid.
  - **Time Zone**: If the Grid member is in a different time zone from the Grid, click **Override** and select a time zone.
  - **Comment**: Type a comment that provides some useful information about the appliance, such as its location.
  - **Master Candidate**: Select this option to designate this appliance as a Master Candidate. For supported vNIOS appliances, see [Supported vNIOS Appliance Models and Specifications](#). Note that the Grid Master Candidate must use the same communication protocol as the Grid Master.
4. Enter the following information about the member that you are adding to the Grid and click **Next**:
  - **Type of Network Connectivity**: Select **IPv6** to configure an IPv6 Grid member or select **IPv4 and IPv6** to configure a dual mode Grid member.
  - Select **Standalone Member** to configure a single member or select **High Availability Pair** to configure an HA member.

For an HA member, enter the **Virtual Router ID** number and if the HA member is configured in dual mode, select **IPv6** in the **Send HA and Grid Communication Over** field.
  - **Required Ports and Addresses**: This table lists the network interfaces based on the type of network connectivity of the Grid member.

For IPv6 Grid member, specify the network information for LAN1(IPv6) interface. For a dual mode Grid member, specify the network information for both LAN1(IPv4) and LAN1(IPv6) interfaces.  
For IPv6 HA member, specify the network information for VIP (IPv6), Node1 LAN1(IPv6), and Node2 LAN1(IPv6) ports. For a dual mode HA member, specify the network information for the following interfaces: VIP (IPv4), Node1 LAN1(IPv4), Node2 LAN1(IPv4), VIP (IPv6), Node1 LAN1(IPv6), and Node2 LAN1(IPv6).

Enter correct information for the following by clicking the field:

    - **Interface**: Displays the name of the interface. You cannot modify this.
    - **Address**: You can enter IPv4 or IPv6 address depending on the type of interface. An IPv6 address is a 128-bit number in colon hexadecimal notation. It consists of eight 16-bit groups of hexadecimal digits separated by colons (example: 2001:db8:0000:0123:4567:89ab:0000:cdef or 2001:db8::123:4567:89ab:0:cdef).
    - **Subnet Mask (IPv4) or Prefix Length (IPv6)**: Enter an appropriate subnet mask for IPv4 address and prefix length for IPv6 address. The prefix length ranges from 2 to 127.
    - **Gateway**: Type the default gateway for the interface. For IPv6 interface, you can also type **Automatic** to enable the appliance to acquire the IPv6 address of the default gateway and the link MTU from router advertisements.
    - **VLAN Tag**: For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
    - **Port Settings**: From the drop-down list, choose the connection speed that you want the port to use. You can also choose the duplex setting. Choose **Full** for concurrent bidirectional data transmission or **Half** for data transmission in one direction at a time. Select **Automatic** to instruct the NIOS appliance to negotiate the optimum port connection type (full or half duplex) and speed with the

connecting switch automatically. This is the default setting. You cannot configure port settings for vNIOS appliances.

- **DSCP Value:** Displays the Grid DSCP value, if configured. To modify, click **Override** and enter the DSCP value. You can enter a value from 0 to 63. For information about DSCP, see [Implementing Quality of Service Using DSCP](#).

5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring DNS for IPv6 Addressing

Configuring the appliance to manage DNS services for IPv6 connections is similar to configuring DNS services for IPv4 connections. For simplicity, the IPv6 procedures are located in the same location as the corresponding procedures for IPv4 in this chapter. In most cases, the key difference within the procedure involves selecting an IPv6 mapping zone instead of an IPv4 mapping zone. You can configure the following tasks:

### *IPv6 DNS Configuration Checklist*

Step	For more information
Create primary or secondary name servers and specify an IPv6 root server.	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Authoritative Zones</a></li> <li>• <a href="#">Specifying a Primary Server</a></li> <li>• <a href="#">Specifying Secondary Servers</a></li> <li>• <a href="#">Creating a Root Zone</a></li> </ul>
Configure the IPv6 zones.	<ul style="list-style-type: none"> <li>• <a href="#">Creating an Authoritative Forward-Mapping Zone</a></li> <li>• <a href="#">Creating an Authoritative Reverse-Mapping Zone</a></li> </ul>
Configure IPv6 resource records	<ul style="list-style-type: none"> <li>• <a href="#">Managing AAAA Records</a></li> <li>• <a href="#">Managing PTR Records</a></li> </ul>

## Configuring DNS Services

This chapter provides general information about DNS service properties and includes the following topics:

- [Configuring DNS Service Properties](#)
- [Clearing DNS Cache](#)
- [Viewing DNS Cache Entries](#)
- [Using Forwarders](#)
- [Controlling DNS Queries](#)
- [Enabling Recursive Queries](#)
- [Controlling AAAA Records for IPv4 Clients](#)
- [About NXDOMAIN Redirection](#)
- [Configuring DNS over TLS and DNS over HTTPS Services](#)
- [Detecting and Mitigating DNS DDoS Attacks](#)
- [Automated Mitigation of Phantom Domain Attacks](#)
- [Detecting NXDOMAIN Attacks](#)
- [Mitigating Possible NXDOMAIN Attacks](#)
- [Support for RRL \(Response Rate Limiting\)](#)
- [Blacklists](#)
- [Configuring Blacklists](#)
- [About Root Name Servers](#)
- [About Sort Lists](#)

- [Configuring a DNS Blackhole List](#)
- [Specifying Hostname Policies](#)
- [About DNS64](#)
- [DNS Record Scavenging](#)
- [Monitoring DNS Queries](#)
- [Configuring DNS Traffic Control Properties](#)

## Configuring DNS Service Properties

You can configure general DNS service properties and change some default values. The DNS service is disabled by default. To enable the member to provide DNS service, you must start the DNS service. For information about how to start and stop the DNS service, see [Starting and Stopping the DNS Service](#). The following topics describe the DNS service properties that you can configure:

- [Configuring DNS Access Control](#)
- [Specifying Time To Live Settings](#)
- [Configuring Hostname and Server ID Options](#)
- [Enabling and Disabling DNS Health Check Monitor](#)
- [Adding an Email Address to the SOA Record](#)
- [Notifying External Secondary Servers](#)
- [Enabling the Configuration of RRset Orders](#)
- [Specifying Port Settings for DNS](#)
- [Using Extension Mechanisms for DNS \(EDNS0\)](#)
- [Deleting PTR Records associated with A or AAAA Records](#)
- [Specifying Minimal Responses](#)
- [Starting and Stopping the DNS Service](#)

## Configuring DNS Access Control

You can add ACEs (access control entries) or use a named ACL (access control list) to determine which hosts can perform specific DNS tasks. For information about how to define a named ACL, see [Defining Named ACLs](#). When you add ACEs or a named ACL to Grid DNS properties, the configuration overrides member and object access control for DNS zone transfers, dynamic DNS updates, DNS queries and recursive queries, blackhole lists, and AAAA filtering. For a full list of operations that support access control, see [Operations that Support Access Control](#).

To configure DNS access control:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the *Grid DNS Properties* editor, click **Toggle Advanced Mode**, and select one of the following tabs for specific DNS tasks:
  - **Updates** tab: Define ACEs or a named ACL to control Grid level dynamic DNS updates, as described in [Enabling DNS Servers to Accept DDNS Updates](#).
  - **Queries** tab: Define ACEs or a named ACL to control Grid level DNS queries, recursive queries, and AAAA filtering, as described in [Controlling DNS Queries](#), [Enabling Recursive Queries](#) and [Controlling AAAA Records for IPv4 Clients](#).
  - **Zone Transfers** tab: Define ACEs or a named ACL to control Grid level DNS zone transfers, as described in [Enabling Zone Transfers](#). This does not apply to zone transfers for Microsoft servers. For information about Microsoft servers, see [Setting Zone Properties](#).
  - **Blackhole** tab: Configure ACEs or a named ACL to define IP addresses and networks that you do not want to include during the DNS resolution process, as described in [Configuring a DNS Blackhole List](#).
  - **DNS64** tab: Configure ACEs or a named ACL for clients to which the appliance sends synthesized AAAA records DNS64 groups, as described in [Setting DNS64 Group Properties](#).
3. Save the configuration.

You can override the Grid settings at the member and object levels.

## Specifying Time To Live Settings

You can specify TTL (time to live) settings for Infoblox host records and resource records. TTL is the time that a name server is allowed to cache data. After the TTL expires, the name server is required to update the data. Setting a high TTL reduces network traffic, but also renders your cached data less current. Conversely, setting a low TTL renders more current cached data, but also increases the traffic on your network.

You can specify global TTL settings at the Grid level, for individual zones, or resource records. When you configure TTL settings for auto-generated records, the following conditions apply:

- NS records that are auto-generated for delegated name servers use TTL settings from their delegated zones.
- Auto-generated glue A and AAAA records use TTL settings from a delegated zone if the name of the name server is below the delegation point and does not belong to an authoritative child zone.
- All other auto-generated NS, A, and AAAA records continue to use TTL settings from their parent zones.
- Auto-generated PTR records do not inherit TTL settings from delegated zones. They use TTL settings from their parent zones.

When you have an RRSET (resource record set) that contains different TTL settings for each record, Grid Manager displays the actual TTL values for these records. However, in DNS responses, the appliance takes the least of the values and returns that as the TTL setting for all resource records in the RRset.

For recursive DNS servers, you can specify the maximum cache TTL value that establishes the time limit for the name server to cache positive responses. You can also specify the maximum negative cache TTL value that specifies the time limit for the name server to cache negative responses. For information about how to configure these settings, see [Specifying Max Cache TTL and Max Negative Cache TTL Settings](#) below.

## Specifying TTL Settings for a Grid

To specify global TTL settings for resource records hosted by Grid members:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the **Basic** tab of the **General** section of the *Grid DNS Properties* editor, modify the following values as necessary:
  - **Refresh**: This interval tells secondary servers how often to send a message to the primary server for a zone to check that their data is current, and retrieve fresh data if it is not. The default is three hours. Ensure that you set the refresh interval to a value above 300 seconds; setting the refresh interval to below 300 seconds may not work as expected.
  - **Retry**: This interval tells secondary servers how long to wait before attempting to recontact the primary server after a connection failure between the two occurs. The default is one hour.
  - **Expire**: If the secondary fails to contact the primary for the specified interval, the secondary stops giving out answers about the zone because the zone data is too old to be useful. The default is 4 weeks.
  - **Default TTL**: Specifies how long name servers can cache the data. The default is eight hours.
  - **Negative-caching TTL (Time to Live)**: Specifies how long name servers can cache negative responses, such as NXDOMAIN responses. The default is 15 minutes.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Specifying TTL Settings for a Zone

To specify TTL settings for host and resource records in a zone:

1. From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab-> *dns\_view* -> *zone* checkbox -> Edit icon.
2. In the *Authoritative Zone* editor, click **Settings**.
3. Click **Override** and complete the fields as described in the preceding section, [Specifying TTL Settings for a Grid](#).

## Specifying the TTL of a Host or Resource Record

To specify the TTL setting for an Infoblox host or resource record:

1. From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns\_view* -> *zone* -> *resource\_record*.
2. The **TTL** tab of the resource record editor displays the TTL setting the resource record inherited from the Grid or zone. Click **Override** and enter a value. The setting is in hours by default. You can change it to seconds, minutes, days or weeks.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

### Specifying TTL Settings for a Lame Server

Servers that are marked as authoritative, but do not respond as authoritative servers are called lame servers. You can specify the number of seconds to cache a lame delegation or lame server indication through the Lame TTL option. Lame TTL usually indicates the amount of time your name server remembers information about the remote name server that is not authoritative for a zone, which is delegated to it.

A domain or sub-domain that is delegated to a server that is not authoritative for the domain is called lame delegation. It indicates that a zone file does not exist for the domain on the server.

The lame time-to-live cache value can be defined at the Grid DNS, Member DNS, or DNS view level. To specify the Lame TTL cache value for a lame delegation or lame server:

1. **Grid:** From the **Grid** tab -> **Grid Manager** tab, select the **DNS** tab, click the **Services** tab -> *member* checkbox, expand the **Toolbar** and click **Edit** -> **Grid DNS Properties**. In the *Grid DNS Properties* editor, select the **General** tab -> click the **Advanced** tab (or click **Toggle Advanced Mode**).  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> *Edit* icon. In the *Member DNS Properties* editor, select the **General** tab -> click the **Advanced** tab (or click **Toggle Advanced Mode**).  
**DNS View:** From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *dns\_view* checkbox -> *Edit* icon. In the *DNS View* editor, select the **General** tab -> click the **Advanced** tab (or click **Toggle Advanced Mode**).
2. In the *Grid DNS Properties*, *Member DNS Properties*, or the *DNS View* editor, select the **General** tab -> click the **Advanced** tab (or click **Toggle Advanced Mode**) and then complete the following:  
**Lame TTL:** Specify the duration of time to cache a lame delegation or lame server. The default value is 600 seconds (ten minutes) and the maximum value is 1800 seconds (thirty minutes). The appliance displays a warning message when you specify a value equal to 0 (zero). The appliance displays an error message when you specify a value greater than 1800 seconds.  
The **Lame TTL** cache value is inherited from the Grid by the member and DNS view levels and this field is disabled, by default. To override the **Lame TTL** cache value, click **Override**. You can override the value at the member and DNS view levels. To retain the same **Lame TTL** value as the Grid, click **Inherit** at the member and DNS view level.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

### Specifying Max Cache TTL and Max Negative Cache TTL Settings

You can specify the maximum duration of time for which your name server caches positive responses using the Max Cache TTL settings. The Max Cache TTL indicates the time limit for which the name server retains records in the cache. When the Max Cache TTL for a record expires, the name server deletes the record from the cache.

You can also specify the maximum duration of time for which your name server caches negative responses through the Max Negative Cache TTL settings. The Max Negative Cache TTL sets the time limit for which the name server retains negative responses (NXDOMAIN/NXRRSET responses) in the cache. The name server deletes a negative response from the cache when the Max Negative Cache TTL period for the entry expires.

You can define the Max Cache TTL value and the Max Negative Cache TTL value at the Grid DNS, Member DNS, and DNS view levels.

To specify the Max Cache TTL and the Max Negative Cache TTL:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the **Toolbar** and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> *Edit* icon.  
**DNS View:** From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns\_view* checkbox -> *Edit* icon.
2. In the *Grid DNS Properties*, *Member DNS Properties*, or the *DNS View* editor, click **Toggle Advanced Mode** if the editor is in the basic mode.
3. Click the **Advanced** subtab of the **General** tab and then complete the following:

- **Max Cache TTL:** Specify the maximum duration of time for which the name server caches positive responses. Select the time period in minutes, hours, or days from the drop-down list. The default value is one week (7 days), and the maximum value is 49710 days, 1193046 hours, or 71582788 minutes. The appliance displays an error message when you enter a value greater than the maximum value. Note that setting the Max Cache TTL value to 0 (zero) will disable the name server from caching any data, and it is not recommended.
  - **Max Negative Cache TTL:** Specify the maximum duration of time for which the name server caches negative responses. Select the time period in minutes, hours, or days from the drop-down list. The default value is three hours, and the maximum value is 7 days, 168 hours, or 10080 minutes. The appliance displays an error message when you enter a value greater than the maximum value. Note that setting the Max Negative Cache TTL value to 0 (zero) will disable the name server from caching negative responses, and it is not recommended.
- The Max Cache TTL value and the Max Negative Cache TTL value are inherited from the Grid at the member and DNS view levels. To override the inherited values, click **Override** and specify the new value. To retain the Grid values, click **Inherit**.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

### Configuring Hostname and Server ID Options

When you configure DNS anycast, multiple DNS name servers share a single IP address. To identify which DNS name server is answering queries, you can configure the hostname and server ID options so the appliance returns the hostname of the DNS name server that is currently answering queries. By default, the hostname and server ID options are disabled on the NIOS appliance. You can configure the hostname bind directive and server-id directive options at the Grid level and override them at the member level. The appliance returns the hostname of the DNS name server that is currently answering queries when a client queries for the hostname.bind or the id.server with record type as TXT and class CHAOS, as follows:

```
dig @<IP> hostname.bind txt CH
```

```
dig @<IP> id.server txt CH
```

To secure the identity of the internet-facing DNS servers, you can configure the hostname and server ID options for specific Grid members that are internet-facing to return a user defined value instead of the real hostname. Alternatively, you can disable the hostname and server ID options at the Grid level and configure them only for those members that are not internet-facing.

To configure the hostname and server ID options:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, and then select **Grid DNS Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.
2. In the *Grid DNS Properties* or the *Member DNS Properties* editor, click **Toggle Advanced Mode** if the editor is in the basic mode.
3. Click the **Advanced** subtab of the **General** tab and then complete the following:
  - **Hostname bind directive:** Select either **Hostname** or **None** from the drop-down list. The default is **None**. If you select **Hostname**, the appliance returns the hostname of the DNS name server that is currently answering queries. Selecting **None** disables the Hostname bind directive option. In the *Member DNS Properties* editor, you can also select **User defined** and specify any hostname of your choice. The appliance returns the specified hostname instead of the real hostname of the DNS name server that is currently answering queries. To override an inherited setting from the Grid, click **Override**. To retain the same setting as the Grid, click **Inherit**.
  - **Server-id directive:** Select either **Hostname** or **None** from the drop-down list. The default is **None**. If you select **Hostname**, the appliance returns the hostname of the DNS name server that is currently answering queries, when a client queries to identify the server ID of the name server that is answering queries. Selecting **None** disables the Server-id directive option. In the *Member DNS Properties* editor, you can also select **User defined** and specify a value of your choice. The appliance returns the specified value when a client queries to identify the server ID of the DNS name server that is answering queries.

To override an inherited setting from the Grid, click **Override**. To retain the same setting as the Grid, click **Inherit**.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Enabling and Disabling DNS Health Check Monitor

You can enable the DNS health check monitor to monitor whether the DNS server is responding to client requests. When you enable this feature, the appliance sends a query to the DNS server and waits for the response until the specified timeout duration. If the appliance is unable to receive a response from the DNS server after the specified number of retries, the appliance sends SNMP traps and email notifications about the failure. The appliance performs the DNS health check periodically based on the specified time interval.



### Warning

*The DNS Health Check monitor might not work properly if DNS blackhole feature is enabled or if any named ACL is blocking the query sent to loopback interface.*

To enable or disable the DNS health check monitor:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, and then select **Grid DNS Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.
2. In the *Grid DNS Properties* or the *Member DNS Properties* editor, click **Toggle Advanced Mode** if the editor is in the basic mode.
3. Click the **Advanced** subtab of the **General** tab and then complete the following:
  - **Enable DNS Health Check:** This checkbox is deselected by default, meaning the DNS health check monitor is disabled. Select this checkbox to enable the DNS health check monitor and specify the following:
    - **Interval:** Enter the time interval in seconds. The interval value is measured from the end of the previous monitoring cycle. The default is 30 seconds. You can enter a value between 10 and 21600 seconds.
    - **Timeout:** Enter the timeout value in seconds. This is the time the appliance waits for a response to the query. The default is 3 seconds. You can enter a value between 1 and 10 seconds.
    - **Retries:** Enter the number of times the appliance tries to send the query after a failed attempt. The default is 3. You can enter a value between 1 and 10.
4. Save the configuration.

## Adding an Email Address to the SOA Record

If the primary name server of a zone is a Grid member, you can add an administrator email address to the SOA record to help admins determine who to contact about this zone.

## Adding an Email Address for SOA Records in the Grid

If all zones hosted by the Grid members have the same administrator, you can add the email address once for the Grid. The appliance then adds the email address to the RNAME field of the SOA records of the zones.

To add an email address to the SOA records at the Grid level:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the **General** -> **Basic** tab of the *Grid DNS Properties* editor, enter the email address in the **E-mail Address (for SOA RNAME field)** field.  
Note the appliance does not support IDN for the **E-mail Address (for SOA RNAME field)** field at the Grid level. You can add an email address containing IDN for the SOA records at the zone level.
3. Save the configuration and click **Restart** if it appears at the top of the screen.



## Adding an Email Address for the Zone SOA Record

To add an email address to the SOA record of a zone:

1. From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab-> *dns\_view* -> *zone* checkbox -> Edit icon.
2. In the *Authoritative Zone* editor, click **Settings**.
3. Click **Override** beside the **Email address (for SOA RNAME field)** field and enter the email address of the zone administrator.
4. Save the configuration and click **Restart** if it appears at the top of the screen.



### Note

The appliance supports IDN for the host name of the **Email address (for SOA RNAME field)** field. For example, you can create [admin@инфоблокс.рф](mailto:admin@инфоблокс.рф) but not [админ@инфоблокс.рф.com](mailto:админ@инфоблокс.рф.com).

## Notifying External Secondary Servers

Grid members can use database replication to maintain up-to-date zone data sets, so the secondary servers in the Grid can keep their zone data synchronized even if the primary server fails. Any external secondary servers can fall out of sync, however, if they rely only on the primary server to send notify messages when there is new zone data. Therefore all authoritative name servers in a Grid (all primary and secondary servers) send notify messages to external secondary servers by default. This ensures that an external secondary name server receives notify messages when its primary server is a secondary name server in a Grid. However, it also increases the number of notify messages.

Infoblox recommends that you do not configure a large number of external secondary servers in stealth mode. To ensure that these secondary servers receive notifications about zone updates, you can allow zone transfers for these IP addresses and then enable the appliance to add them to the also-notify statement. For information about how to configure this feature, see [Configuring Zone Transfers](#).

To specify whether secondary name servers in the Grid are to send notify messages to external secondary name servers:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the *Grid DNS Properties* editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Advanced** subtab of the **General** tab.
4. Complete the following:
  - **Enable Grid secondaries to notify external secondaries:** This option is enabled by default.
  - **Notify Delay:** Specify the number of seconds that the Grid secondary servers delays sending notification messages to the external secondaries. The default is five seconds.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

For the external secondary servers to accept notify messages from the secondary name servers in the Grid and then request zone transfers from them, you must configure the external secondary servers to use the Grid secondary servers as the source of the zone transfers. This ensures that the external secondary servers continue to receive notify messages, even if the primary server is unavailable.

## Enabling the Configuration of RRset Orders

You can use the Infoblox GUI to configure the order that the appliance uses to return the A and AAAA records associated with an Infoblox host. This feature is useful when you want the appliance to return the A and AAAA records of a host in a specific order. For example, if you want the management address to appear first on a list of IP addresses associated with a network device, you can configure the order of the IP addresses so the management address is always returned first on the list when you look up the name of the device. For information about using the Infoblox API to configure RRset order (resource record order) of a host, refer to the *Infoblox API Documentation*.

To specify an RRset order of a host record, you must first enable the feature at the Grid level. When you enable this feature and there are multiple IP addresses associated with the host record, you can specify one of the following RRset orders through the *Host Record* wizard and editor:



- **Fixed:** The A and AAAA records of the host are returned in the order that you specify in the IPv4 and IPv6 address tables.
- **Random:** The A and AAAA records of the host are returned in a random order.
- **Cyclic:** The A and AAAA records are returned in a round robin pattern.

For information about specifying RRset order of a host record, see [Adding Host Records](#).

Note that when you configure an order type for the IP addresses associated with a host record, the order type applies to both the A and AAAA records of the host. It does not apply to any non-host A or AAAA records that may have the same owner name as the host record. By default, the appliance returns resource records in a cyclic or round robin order. The return order of non-authoritative data retrieved from a recursion is not affected by the host RRset order, and that remains cyclic.

When you enable the RRset order for hosts at the Grid level, you may not be able to maintain the same DNS responses on a recursive server if it exists in the Grid. You can preserve the original cached DNS responses by configuring a fixed RRset order on the recursive server so it can return A and AAAA records associated with domain names in the original order they were received. For information about configuring the RRset order for the cached DNS responses, see [Preserving the RRset Order for Cached DNS Responses](#) below.

To enable the configuration of RRset order for a host record:

1. From the **Data Management** tab -> **DNS** tab, expand the Toolbar, and then click **Grid DNS Properties**.
2. In the *Grid DNS Properties* editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Advanced** subtab of the **General** tab.
4. Complete the following:
  - **Enable setting RRset order for hosts with multiple addresses:** Select this checkbox to enable the configuration of RRset order for a host record. After you enable this feature, you can configure the RRset order in the *Host Record* wizard or editor. For information, see [Adding Host Records](#).
  - **Preserve host RRset order for Grid secondaries that use DNS zone transfers:** This is enabled only when you have enabled the setting of RRset order for host records. When you select this checkbox, the RRset order that you configure for a host record applies to the resource records of the Grid secondaries that are in the DNS transfer mode.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

### Preserving the RRset Order for Cached DNS Responses

By default, when a client queries a domain name, the DNS caching appliance returns the A and AAAA records of the domain name in the cyclic order. However, this default behavior can be overridden if you have enabled and configured (at the Grid level) fixed RRset order for hosts that have multiple addresses. When you override the default behavior and preserve the fixed RRset order for cached DNS responses, the DNS caching appliance returns A and AAAA records associated with domain names in the order they were received from an upstream server. You can preserve the RRset order for the cached DNS responses and specify the fixed RRset order for A, AAAA, or both A and AAAA records at the Grid level and override at the member and DNS view levels. Note that configuring fixed RRset order for specific FQDNs might slightly affect the performance of the DNS caching appliance.

To preserve the fixed RRset order for cached DNS responses at the Grid, member, or DNS view level:

1. **Grid:** From the **Data Management** tab -> **DNS** tab, expand the Toolbar, and then click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
**DNS View:** From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *dns\_view* checkbox -> Edit icon.
2. In the editor, select the **RRset Order** tab -> click the **Basic** tab, and then complete the following:
  - **Enable fixed RRset order for following FQDNs:** Select this checkbox to preserve the configuration of RRset order for cached DNS responses.
  - In the FQDN table, specify the list of FQDN entries for which you want to preserve the RRset order. Note that you can configure a maximum of 25 FQDNs for the specified RRset order.  
 You can click the Add icon and complete the following to add a new entry to the list:
    - **FQDN:** Enter the fully qualified domain name with which the A or AAAA record is associated.
    - **Record Type:** Select the record type from the drop-down list. You can select **A**, **AAAA**, or **Both A and AAAA**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Specifying Port Settings for DNS

When requesting zone transfers from the primary server, some secondary DNS servers use the source port number (the primary server used to send the notify message) as the destination port number in the zone transfer request. If the primary server uses a random source port number when sending the notify message—that the secondary server then uses as the destination port number when requesting a zone transfer—zone transfers can fail if there is an intervening firewall blocking traffic to the destination port number.

Specifying a source port number for recursive queries ensures that a firewall allows the response. If you do not specify a source port number, the NIOS appliance sends these messages from a random port number. You can also specify a source for the DNS Traffic Control health check.

When performing recursive queries, the NIOS appliance uses a random source port number above 1024 by default. The queried server responds using the source port number in the query as the destination port number in its response. If there is an intervening firewall that does not perform stateful inspection and blocks incoming traffic to the destination port number, the recursive query fails.

You can specify a source port number for notify messages to ensure the firewall allows the zone transfer request from the secondary server to the primary server. If you do not specify a source port number, the NIOS appliance sends messages from a random port number above 1024.

You can limit If you have configured anycast and non-anycast IP addresses on the loopback interface, you must enable the appliance to provide DNS services on them. You can also configure the appliance to listen for DNS queries on a specific IP address that you configure on the loopback interface, by separating the source port for DNS queries from the port for notify messages and zone transfer requests. For information about the loopback interface and anycast addressing, see [Configuring IP Addresses on the Loopback Interface](#).

You can specify the source address settings for a Grid member and for DNS views assigned to a Grid member. Note that you can specify the source address settings for only specific DNS views that contain zones that are assigned to a Grid member. The static source port values for DNS views are inherited from the Member DNS properties.

## Specifying Source Ports

To specify port numbers and settings for queries, notify messages and zone transfer requests for a Grid member or DNS view assigned to a Grid member:

1. **Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
**DNS view:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.
2. In the *Member DNS Properties* editor, click **General** -> **Basic** tab.
3. You can change the port settings as follows:
  - **Listen on these additional IP addresses:** Click the Add icon to add an anycast or non-anycast address you configure on the loopback or VLAN interface. You must add all IP addresses you configure on the loopback or VLAN interface so the appliance can provide DNS services on them. Adding source ports for listening supports both IPv4 and IPv6 interfaces. For information about adding IP addresses on the loopback interface, see [Configuring IP Addresses on the Loopback Interface](#).  
Enter correct information for the following by clicking the field in the row:
    - **Address:** Type the IPv4 or IPv6 address depending on the type of interface.
    - **Interface:** The Interface column displays the name of the interface.
    - **VLAN Tag:** For a VLAN, enter the VLAN tag or ID. You can enter a number from 1 to 4094. Ensure that you configure the corresponding switch accordingly.
  - **Send queries from:** If you want to improve the DNS service performance, you can separate the DNS queries from the notify messages and zone transfer requests. Select a value from the drop-down list to select an interface name: **VIP, MGMT, LAN2, Any interface, IP**.
    - **IP Address:** This is displayed only when you select **IP** from the drop-down list. Specify the IP address of the source.
  - **Send notify messages and zone transfer requests from:** From the drop-down list, select the source port of the notify messages and zone transfer requests that the Grid member sends. Select a value from the drop-down list to select an interface name: **VIP, MGMT, LAN2, Any interface, IP**. You can select IP addresses on the loopback or non-primary VLAN interface.

- **IP Address:** This is displayed only when you select **IP** from the drop-down list. Specify the IP address of the source.

Note:

- **IP** is displayed only if you have additional IP addresses such as the loopback address or VLANs configured.
- If you select IP addresses on the loopback or non-primary VLAN interface, then you must add these IP addresses in the **Listen on these additional IP addresses** table.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

If you have enabled secondary name servers in the Grid to send notify messages to external secondary name servers, you can specify the delay time for sending the notify messages. For information about enabling Grid secondary servers to send notification messages to the external secondaries, see [Notifying External Secondary Servers](#).

To specify the delay time for the Grid secondary servers to send notify messages to the external secondaries:

1. **Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
**DNS view:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.
2. In the editor, click **Toggle Advanced Mode**.
3. **Member:** When the additional tabs appear, click the **Advanced** subtab of the **General** tab.  
**DNS view:** When the additional tabs appear, click the **Advanced** subtab of the **DNS Views** tab.
4. Complete the following:
  - **Notify Delay:** Specify the number of seconds that the Grid secondary servers delays sending notification messages to the external secondaries. You can enter a value between 5 and 86400 seconds. The default is five seconds.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

### Specifying Static Source Ports

To specify static source ports for a Grid and Grid member:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
**DNS View:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.
2. In the editor, click **Toggle Advanced Mode**.
3. **Member:** When the additional tabs appear, click the **Advanced** subtab of the **General** tab. **DNS View:** When the additional tabs appear, click the **Advanced** subtab of the **DNS Views** tab. To override an inherited property, click **Override** next to it and complete the appropriate fields.
4. Complete the following:
  - **Set static source UDP port for queries (not recommended):** This is disabled by default. To override the value that has been inherited from the Grid, click **Override**. Select this checkbox to enable it and enter the UDP port number. You can enter a value between 1 and 63999. To retain the same value as the Grid, click **Inherit**.
  - **Set static source UDP port for notify messages:** This is disabled by default. To override the value that has been inherited from the Grid, click **Override**. Select this checkbox to specify a source port for notify messages to ensure that the firewall allows the zone transfer request from the secondary server to the primary server. You can enter a value between 1 and 63999. If you do not specify a source port, the appliance sends messages from a random port with a number above 1024. To retain the same value as the Grid, click **Inherit**.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

### Using Extension Mechanisms for DNS (EDNS0)

The NIOS appliance supports EDNS0 (Extension Mechanisms for DNS), which allows DNS clients to expand and advertise up to 4096 bytes of UDP packets for certain DNS parameters. EDNS0 facilitates the transfer of UDP packets beyond the original restricted packet size of 512 bytes. As defined in RFC 6891, EDNS0 provides extended UDP packet size that supports additional DNS functionality, such as DNSSEC. When EDNS0 is supported, the DNS client adds

information to the additional data section of a DNS request in the form of an OPT pseudo-RR (resource record). An OPT RR does not contain actual DNS data and its contents pertain to the UDP transport layer message only. An OPT RR is not cached, forwarded, or stored. For more information about EDNS0, refer to RFC 6891 Extension Mechanisms for DNS (EDNS0).

By default, EDNS0 is enabled on the NIOS appliance and all outgoing recursive queries are set to have a maximum UDP packet size of 1220 bytes. The EDNS0 UDP packet size (EDNS0 buffer size) is configurable and can be set from a minimum of 512 bytes to a maximum of 4096 bytes. Typically, when the appliance receives a DNS request that contains an OPT RR, it assumes the DNS client supports EDNS0 and thus scales its response accordingly. When the appliance is used as a forwarder or a resolver for recursive queries and communicates with a client that does not support EDNS0, the appliance sends three queries starting with one that contains EDNS0 and DNSSEC support messages and is set to a maximum UDP packet size of 4096 bytes. When the first query fails, the appliance sends another query that contains only the EDNS0 support message. If the second attempt fails too, the appliance sends a third query that indicates a standard 512-byte query. Note that when EDNS0 is not used, DNS packets may be sent over TCP. For DNS service to function properly at this stage, ensure that you configure your firewall accordingly.

The following information demonstrates how the appliance responds when EDNS0 is enabled by default and the end server does not support EDNS0:

```
Packet 0954: 08:19:38.925 - query for www.google.com from Infoblox to forwarder (with EDNS0 support by setting the Extended Label Type to '01' and DNSSEC OK bit to '1')
```

```
Packet 1138: 08:19:47.927 - query for www.google.com from Infoblox to forwarder (with EDNS0 support by setting the Extended Label Type to '01' and DNSSEC OK bit to '0')
```

```
Packet 1504: 08:19:58.929 - query for www.google.com from Infoblox to forwarder (without EDNS0 and DNSSEC support by sending a standard 512-byte query)
```

```
Packet 1505: 08:19:30.960 - query response for www.google.com from forwarder to Infoblox
```

### Configuring the EDNS0 Buffer Size and UDP Buffer Size

When the size of a DNS message that is transferred over UDP exceeds the maximum transmission unit, IP packets get fragmented and reassembled. IP fragmentation is considered fragile as it makes certain attacks on DNS possible. NIOS allows you to configure the EDNS0 buffer size and UDP buffer size attributes to control the data packet size allowed in DNS responses so that the data is transferred without fragmentation.

- The UDP buffer size is used by authoritative DNS servers when data is transferred between DNS server and DNS client to ensure that DNS messages they send are not larger than the configured buffer size.
- The EDNS0 buffer size is used by recursive DNS servers when data is transferred between DNS servers to ensure that DNS messages they send are not larger than the configured buffer size.



#### Note

The UDP buffer size and EDNS0 buffer size attributes are available only for BIND resolvers and not for unbound resolvers.

You can configure the EDNS0 buffer size and the UDP buffer size are configurable for a Grid, member, standalone system, and a DNS view. To configure the EDNS0 buffer size and UDP buffer size, complete the following steps:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab, and then click the **Members** tab -> *member* checkbox -> Edit icon.

To override an inherited property, click **Override** next to it and complete the appropriate fields.

**Standalone system:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **System DNS Properties**.

**DNS view:** From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *dns\_view* checkbox -> Edit icon. Note that if there is only one DNS view, for example, the predefined default view, then you can just click the Edit icon beside it.

To override an inherited property, click **Override** next to it and complete the appropriate fields.

2. In the *Grid DNS Properties*, *Member DNS Properties*, *System DNS Properties*, or *DNS View* editor, click **Toggle Advanced Mode** if the editor is in basic mode.
3. Click the **General** tab -> **Advanced** tab, and complete the following:
  - a. **EDNS0 Buffer Size:** Specify the maximum packet size to be allowed in DNS query responses when transferring DNS messages between DNS servers. The default buffer size is 1220 bytes. The minimum buffer size that you can set is 512 bytes and the maximum is 4096 bytes. Infoblox recommends that you configure the **EDNS0 Buffer Size** value in the range of 512 to 1220 bytes. DNS responses that exceed 1220 bytes can get fragmented and may result in unexpected behavior when resolving queries. Note that if you want DNS query responses to use the configured EDNS0 buffer size in servers that support EDNS0, then ensure that you do not disable EDNS0.
  - b. **UDP Buffer Size:** Specify the maximum packet size to be allowed in DNS query responses when transferring DNS messages from DNS servers to DNS clients. The default buffer size is 1220 bytes. The minimum buffer size that you can set is 512 bytes and the maximum is 4096 bytes. Infoblox recommends that you configure the **UDP Buffer Size** value in the range of 512 to 1220 bytes. DNS responses that exceed 1220 bytes can get fragmented and may result in unexpected behavior when resolving queries.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Disabling EDNS0

To ensure that end servers that do not support EDNS0 can respond to recursive queries from the NIOS appliance and to improve DNS performance, you can disable EDNS0 for the Grid and override the Grid settings for individual members. Note that you cannot configure the maximum UDP packet size, which is set for 4096 bytes by default. When you disable EDNS0, the appliance does not include OPT RRs for all outgoing recursive DNS queries. Thus remote end servers that do not support EDNS0 can still respond to the queries. This feature is useful when your NIOS appliance is used as a forwarder or a resolver for recursive queries, and the end servers in the configuration do not support EDNS0.



### Warning

*When you disable EDNS0, all outgoing DNSSEC queries to zones within trusted anchors will fail even if DNSSEC validation is enabled. This is due to the restriction of the UDP packet length when you disable EDNS0. For more information about DNSSEC, see [Configuring DNSSEC](#).*

To disable EDNS0:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click the **General** tab -> **Advanced** tab, and complete the following:
  - **Disable EDNS0:** This checkbox is deselected and EDNS0 is enabled by default. To override the value inherited from the Grid, click **Override**. To retain the same value as the Grid, click **Inherit**. Select this checkbox to disable EDNS0. When you disable EDNS0, the appliance does not include OPT RRs for all outgoing recursive DNS queries and all outgoing DNSSEC queries to zones within trusted anchors will fail even if DNSSEC validation is enabled.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Deleting PTR Records associated with A or AAAA Records

You can configure the appliance to confirm whether to delete a PTR record when its corresponding A or AAAA record is being deleted. This feature is valid only if you have configured the appliance to automatically generate a PTR record when you create an A or AAAA record. For information about adding an A or AAAA record, see [Adding A Records](#) and [Adding AAAA Records](#). When you delete a resource record, the appliance moves it to the Recycle Bin, if enabled. You can later restore it if needed. Note that this option is disabled by default for all new installations.

To enable this option:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the *Grid DNS Properties*, click the **General** tab -> **Advanced** tab, and complete the following:
  - **Enable PTR record removal for A/AAAA records:** This checkbox is deselected by default. When you select this checkbox, the appliance displays the *Delete Confirmation* dialog box to confirm that you want to delete a PTR record associated with the A or AAAA record that is being deleted. For information about deleting an A or AAAA record, see [Modifying, Disabling, and Deleting Host and Resource Records](#).
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Specifying Minimal Responses

A NIOS appliance returns a minimal amount of data in response to a query, by default. It includes records in the authority and additional data sections of its response only when required, such as in negative responses. This feature speeds up the DNS services provided by the appliance.

To disable returning minimal responses:

1. From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.
2. In the *Member DNS Configuration* editor, click **General** -> **Basic** tab.
3. Clear the Return minimal responses checkbox.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Starting and Stopping the DNS Service

The DNS service is disabled by default. After you complete the DNS configuration, you can start DNS service on a member. You can also disable the DNS service on any Grid member. Be aware that disabling the DNS service on a member removes the NS records from it. If you later re-enable DNS service for this member, the NS records are then restored.

To start DNS service on a member:

1. From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox -> click the Edit icon.
2. In the *Member DNS Properties* editor, click **General** -> **Basic** tab, and do the following:
  - **LAN1:** Select **IPv4** to start the DNS service using IPv4 or select **IPv6** to start the DNS service using IPv6. Note that for a dual mode member, **IPv4** will be selected by default.
3. Save the configuration.
4. Expand the Toolbar and click **Start**.
5. In the *Start Member DNS Service* dialog box, click **Yes**.  
Grid Manager starts the DNS service on the selected member.

You can stop DNS service on a member by selecting the member checkbox and click **Stop** from the Toolbar.

## Clearing DNS Cache

The NIOS appliance allows you to clear certain information from the DNS cache. You can do the following:

- Clearing DNS Cache
- Clearing Cache for DNS Views
- Clearing Domains and Subdomains from DNS Cache



## Clearing DNS Cache

You can clear all the entries that are saved in the DNS cache. When you clear DNS cache on the NIOS appliance, entire BIND recursive cache is cleared.

To clear DNS cache:

1. From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox.
2. Expand the Toolbar, click **Clear** -> **Clear DNS Cache**.
3. Click **Yes** in the confirmation dialog box to clear DNS cache.

## Clearing Cache for DNS Views

You can configure the NIOS appliance to clear cache of a specific DNS view. This feature clears cache entries of a specific DNS view that is associated with the selected member.

To clear cache of a DNS view:

1. From the **Data Management** tab, select the **DNS** tab -> click the **Members** tab.
2. Expand the Toolbar, click **Clear** -> **Clear View's Cache**.
3. Specify the following in the *Clear View's Cache* dialog box:
  - **Member:** Click **Select Member** to select a member. If there are multiple members, the *Member Selector* dialog box is displayed, from which you can select a member. Click the required member name in the dialog box. You can also click **Clear** to clear the displayed member and select a new one.
  - **DNS View:** Select a **DNS View** from the drop-down list. This list box appears only when there are multiple DNS views in the network view.
  - Click **Clear Cache** to clear the cache entries of the corresponding DNS View.



### Note

The entire name server recursive cache is cleared, if you do not specify a DNS view when you clear cache using **Clear View's Cache** and **Clear Domain Name** features on the NIOS appliance.

## Clearing Domains and Subdomains from DNS Cache

You can clear a specific domain and its subdomains from the DNS cache. When you clear a domain from the DNS cache, it is also removed from the BIND recursive cache.

To clear a specific domain and its subdomains:

1. From the **Data Management** tab, select the **DNS** tab -> click the **Members** tab.
2. Expand the Toolbar, click **Clear** -> **Clear Domain Name**.
3. Specify the following in the *Clear Domain Name from Cache* dialog box:
  - **Domain Name:** Enter a domain name you want to delete.
  - **Clear entire domain (including subdomains):** Select the checkbox to clear the specified domain and its subdomains from the DNS cache. For example, if you enter `corpxyz.com` in the **Domain Name** field, then selecting this checkbox clears the entire domain including its subdomains such as `www.corpxyz.com`, `corpxyz.com`, `x.corpxyz.com`, etc. This checkbox is deselected by default.
  - **Member:** Click **Select Member** to select a member. If there are multiple members, the *Member Selector* dialog box is displayed, from which you can select a member. Click the required member name in the dialog box. You can also click **Clear** to clear the displayed member and select a new one.
  - **DNS View:** Select a **DNS View** from the drop-down list. This list box appears only when there are multiple DNS views in the network view.
  - Click **Clear Domain Name** to clear the domain name from the cache.

## Viewing DNS Cache Entries

The NIOS appliance allows you to view certain information that is stored in the DNS cache. You can do the following:

- View DNS configuration
- View DNS cache details
- View statistics

## Viewing DNS Configuration

The NIOS appliance enables you to view DNS configuration details. You can view the configuration details by using a browser. To view DNS configuration:

1. From the **Data Management** tab, select the **DNS** tab -> click the **Members** tab.
2. Expand the Toolbar, click **View** -> **View Configuration**.

## Sample Output

```
include "/infoblox/var/named_conf/tsig.key";
options {
zone-statistics yes;
directory "/infoblox/var/named_conf"; version none;
hostname none; recursion yes;
listen-on { <loopback address>; 10.34.1.18; };
query-source address 10.34.1.18 port *; notify-source 10.34.1.18 port *;
transfer-source 10.34.1.18;
minimal-responses yes; max-cache-size 536870912; infoblox-top-query yes;
infoblox-top-query-log-interval 60;
infoblox-top-query-client 500;
infoblox-top-query-name 500;
infoblox-top-query-rr-type 500;
infoblox-top-query-nxdomain 500;
infoblox-top-query-servfail 500;
infoblox-top-query-rpz 99;
infoblox-top-query-rpz-items-per-client 100;
lame-ttl 600;

# for service restart: allow_bulkhost_ddns = Refusal allow-transfer { any; };
forwarders { 10.32.0.177; };
avoid-v4-udp-ports { 2114; 2113; 2115; 8000; 8089; 9997; 2222; 7911; 7912;
8000; 8089; 9997;
8080; 9000; 9999; 9004; 2022; 3374; 3115; 1194; };
transfer-format many-answers;
};

# Worker threads: default
# Bulk Host Name Templates:
```



```

#Four Octets: "$1$2-$3-$4" (Default)
#One Octet: "-$4"
#Three Octets: "$2$3-$4"
#Two Octets: "$3$4"
include "/infoblox/var/named_conf/dhcp_updater.key";
include "/infoblox/var/named_conf/rndc.key";
controls {
inet <loopback address> port 953
allow { <loopback address>; } keys { "rndc-key"; };
};
logging {
channel ib_syslog { syslog daemon; severity info;
};
category default { ib_syslog; }; category rpz { null; };
};
acl all_dns_views_updater_keys { key DHCP_UPDATER_default; key DHCP_UPDATER1;
key DHCP_UPDATER3; };

```

#### Corresponding Dashboards and Settings

- The `infoblox-top-query-log-interval 60` setting logs the DNS top queries with an interval of 60 seconds.
- The `infoblox-top-query-name 500` setting captures the DNS top query domain names with the limit of 500.
- The `infoblox-top-query-client 500` setting corresponds to the *DNS Domain Queried by Client* dashboard. For more information about the dashboard, see *Predefined Dashboards*.
- The `infoblox-top-query-rpz` and `infoblox-top-query-rpz-items-per-client` settings are RPZ-specific settings. They correspond to the **Top N Limit** and **Total RPZ Entries per Client** fields under *Grid Reporting Properties > DNS > DNS RPZ Rule Hit Configuration*. For more information, see *Grid Reporting Properties*.
- You can also choose the number of filter options to display by using the Top N . The default is 10. For more information, see *About Dashboards*.

#### Viewing DNS Cache Details

You can view data stored in cache for the DNS views that are configured in the NIOS appliance. You can view the details through a browser.

To view cache details:

1. From the **Data Management** tab, select the **DNS** tab -> click the **Members** tab.
2. Expand the Toolbar, click **View** -> **View Cache**.

## Sample Output

```
;; Start view _default
;;; Cache dump of view '_default' (cache _default)
;$DATE 20121018180555
; authanswer a.test.com .23876IN A4.4.4.4
;; Address database dump
;; Dump complete
```

## Viewing Statistics

The View Statistics feature enables you to view DNS Statistics of a Grid member. You can view statistics through a browser.

To view statistics:

1. From the **Data Management** tab, select the **DNS** tab -> click the **Members** tab.
2. Expand the Toolbar, click **View** -> **View Cache**.

You can view statistics in the *DNS Statistics for Member* dialog box.

## Using Forwarders

A forwarder is essentially a name server to which all other name servers first send queries that they cannot resolve locally. The forwarder then sends these queries to DNS servers that are external to the network, avoiding the need for the other name servers in your network to send queries off-site. A forwarder eventually builds up a cache of information, which it uses to resolve queries. This reduces Internet traffic over the network and decreases the response time to DNS clients. This is useful in organizations that need to minimize off-site traffic, such as a remote office with a slow connection to a company's network.

You can select any Grid member to function as a forwarder. You must configure your firewall to allow that Grid member to communicate with external DNS servers. You can also configure NIOS to send queries to one or more forwarders. You can define a list of forwarders for the entire Grid, for each Grid member, or for each DNS view.

If your network configuration includes Infoblox BloxOne Threat Defense, you can configure NIOS Grid members (physical or virtual appliance) to forward recursive queries to BloxOne Threat Defense. For more information about BloxOne Threat Defense, see [BloxOne Threat Defense](#). For information about how to configure NIOS members as DNS forwarding proxies, see [Forwarding Recursive Queries to BloxOne Threat Defense](#) below.

## Selecting Forwarders

When there is more than one forwarder in the Grid, the NIOS resolver uses a smoothed metric derived from RTT (Round Trip Time) to select the name server to send queries to. RTT is the length of time between when a query was sent and when its response was received.

## Specifying Forwarders

To configure forwarders for a Grid, member, or DNS view, complete the following:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox -> Edit icon.  
**DNS View:** From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *dns\_view* checkbox -> Edit icon.  
Note that if there is only one DNS view—for example, the predefined default view—you can just click the Edit icon beside it.  
To override an inherited property, select **Override** next to it and complete the appropriate fields.

2. Click the **Forwarders** tab.
3. Click the Add icon.
4. Enter an IP address in the text field. The field supports entry for both IPv4 and IPv6 values.
  - a. To remove a forwarder, select the IP address from the Forwarders list, and then click the Delete icon.
  - b. To move a forwarder up or down on the list, select it and click the **Up** or **Down** arrow.
5. To use only forwarders on your network (and not root servers), select the **Use Forwarders Only** checkbox.
6. Select the **Add client IP, MAC addresses, and DNS View name to outgoing recursive queries** checkbox to include the client IP address, MAC address, and DNS view name of the client from which the DNS query was initiated, to outgoing recursive queries. For information on recursive queries, see [Enabling Recursive Queries](#). Selecting this option includes EDNS0 custom options.
7. Select the **Copy client IP, MAC addresses, and DNS View name to outgoing recursive queries** checkbox to copy and validate the client IP address, MAC address, DNS view name from incoming queries to outgoing queries. If this checkbox is selected and:
  - Only one custom option is present, the IP address or MAC address or DNS view name is copied to the outgoing query without adding the missing option. An incoming query can contain only one IP address or MAC address or DNS view name.
  - No custom option is present, if the **Add client IP, MAC addresses, and DNS View name to outgoing recursive queries** checkbox is selected, valid IP address, MAC address, and DNS view name EDNS0 options are copied from incoming queries to outgoing recursive queries without any change. If the **Add client IP, MAC addresses, and DNS View name to outgoing recursive queries** checkbox is not selected, no options are added to outgoing recursive queries.
 For more information about EDNS0 options, see [Configuring DNS Traffic Control Properties](#) and [Using Extension Mechanisms for DNS \(EDNS0\)](#).
8. Save the configuration and click **Restart** if it appears at the top of the screen.



#### Note

Infoblox recommends that you do not include client IP addresses and MAC addresses in queries directed to non-Infoblox DNS servers and that you include the addresses in only those queries directed at Infoblox DNS servers.

#### Forwarder Limitations

- Forwarding zones (also known as conditional forwarders) do not support the **Add client IP, MAC addresses, and DNS View name to outgoing recursive queries** and the **Copy client IP, MAC addresses, and DNS View name to outgoing recursive queries** checkboxes.
- Only BIND-based DNS servers support these options. Unbound-based DNS servers do not support these options.

#### Forwarding Recursive Queries to BloxOne Threat Defense

To forward recursive queries to BloxOne Threat Defense, you must first register each NIOS member in your Grid as a DNS forwarding proxy through the Cloud Services Portal. When you register a Grid member, the DNS forwarding proxy software is installed on the member. The DNS forwarding proxy embeds the client IP addresses in the DNS queries before forwarding them to BloxOne Threat Defense. The communications are encrypted and client visibility is maintained. Once you set up a DNS forwarding proxy on a Grid member, all recursive queries for that member are forwarded to a local DNS forwarding proxy by the NIOS DNS service. It also caches responses to speed up DNS resolution for future queries. For information about configuring DNS forwarding proxies, see [On-Prem Host Management](#).

Make sure that port 443 is open against its respective domain for DNS forwarding proxy to work between NIOS and BloxOne Threat Defense.

#### Guidelines When Enabling Recursive Query Forwarding on a Grid Member

Note the following when you enable recursive query forwarding on a Grid member:

- Make sure that you enable recursion on the member that you wish to use as a forwarding proxy to BloxOne Threat Defense. For information about how to enable recursion on a Grid member, see [Enabling Recursive Queries](#).
- DNS forwarding proxy does not work on systems configured in the IPv6-only mode.
- DNS forwarding proxy is not supported on the IB-100, IB-810, IB-820, IB-V810, and IB-V820 appliances. Infoblox recommends that you do not configure DNS forwarding proxy on these appliances.
- DNS forwarding proxy is not supported on any appliance that is running on a memory lower than 4 GB.
- Grid Manager ignores global forwarders and all recursive queries are sent to BloxOne Threat Defense.
- Unbound is not supported on a Grid member when it uses Bind to send recursive queries to BloxOne Threat Defense. For information about Unbound, refer to the *Infoblox DNS Cache Acceleration Application Guide*.
- There might be a significant performance impact on your appliance and network during the DNS forwarding proxy installation process depending on the network connectivity between NIOS and BloxOne Threat Defense. Every node will have to install the DNS forwarding proxy before serving DNS recursive queries, which includes the HA nodes.
- When you enable DNS forwarding to BloxOne Threat Defense, the QPS (query per second) throughout might vary, depending on your appliance models and the cache hit ratios. You might see a bigger performance impact when the cache hit ratio is lower.
- DNS forwarding proxy does not work with DNSSEC in case a request was redirected by BloxOne Threat Defense. If you are running DNS forwarding proxy on NIOS, you must disable DNSSEC validation. Even if you disable DNSSEC validation, validation still takes place through BloxOne Threat Defense. To enable DNS forwarding proxy to work with DNSSEC in case a request was redirected by BloxOne Threat Defense, see [Enabling DNS Forwarding Proxy to Work with DNSSEC](#) below.
- To bypass recursive query forwarding to BloxOne Threat Defense, you must disable the DNS forwarding proxy service.
- By adding the join token you obtained from the Cloud Services Portal, and specifying the IP address of the Cloud Services Portal in the **CSP Resolver** field, you can establish connectivity between NIOS and Cloud Services Portal. Thereafter, you can enable the NIOS Grid Connector service in the Cloud Services Portal. This capability provides you with a single interface for viewing comprehensive network data such as global IP space, subnets, IP addresses, and DHCP lease data for your BloxOne Cloud infrastructure and NIOS. For more information, see [Configuring NIOS Grid Connector](#) in the BloxOne Threat Defense documentation.

## Enabling a Grid Member to Forward Recursive Queries to BloxOne Threat Defense

To enable a Grid member to forward recursive queries to BloxOne Threat Defense, complete the following:

1. Log in to the Cloud Services Portal at [csp.infoblox.com](https://csp.infoblox.com).
2. Create a join token by following the instructions in the [Managing Join Tokens for On-Prem Hosts](#) section of the BloxOne Threat Defense documentation. In an HA environment, two on-prem hosts are created. You must ensure that the configurations for both these on-prem hosts are the same for the HA nodes to work seamlessly.
3. Log in to Grid Manager.
4. On the **Grid** tab, click the **Grid Manager** tab -> **Grid Properties** -> **Edit**.
5. In *Grid Properties Editor*, click the **CSP Config** tab and complete the following:
  - **Join Token:** Configure the join token that you created in the Cloud Services Portal in step 2. However, if the field is empty, the cloud connection is not to be terminated.
  - **CSP Resolver:** Displays the IP address of the local DNS resolver. This IP address or DNS is used to resolve Infoblox domains when the DNS Forwarding Proxy service starts. You must configure at least one external resolver that will be used to resolve all required domains. If you do not enter an IP address, 52.119.40.100 is taken as the default.
  - **HTTP proxy:** Enter the URL of the proxy server in the `http://<IP/host>:<port>` format. When you update the HTTP proxy, the NIOS on-prem agent updates it to the other on-prem containers by restarting the containers at a specific interval which can cause a maximum delay of 15 minutes.
6. Click **Save & Close**.
7. **Member:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox -> Edit icon.
8. In the *Grid Member Properties Editor*, click the **CSP Config** tab, and then complete the following. To override an inherited property, click **Override** next to it and enter the value for the appropriate fields if you do not want to inherit the values from the Grid. Once you override, the settings are applicable only at the member level.
  - a. **Join Token:** Displays the join token value that is inherited from the Grid. However, if the field is empty, the cloud connection is not to be terminated.

- b. **CSP Resolver:** Displays the Cloud Services Portal resolver value that is inherited from the Grid.
  - c. **HTTP Proxy:** Displays the URL that is inherited from the Grid.
  - d. **Standalone:** Select this option when the member is standalone.
    - i. **Access Key:** You cannot edit the value of this field; you can only clear it. However, clearing the access key value does not terminate the cloud connection.
  - e. **HA Enabled:** Select this option when the member is an HA.
    - i. **Access Key:** You cannot edit the value of this field; you can only clear it. In case of a NIOS upgrade, the access keys are the same for both the active and passive nodes.
9. Click **Save & Close**.
  10. On the **Grid** tab, select the **Grid Manager** tab -> **DFP** tab -> *member* checkbox -> Edit icon.
  11. In *Member DFP Properties* editor, select the **Fallback to the default resolution process if BloxOne Threat Defense Cloud does not respond** checkbox to forward recursive queries to the local root name servers in case BloxOne Threat Defense fails or if BloxOne Threat Defense fails to resolve recursive queries. For newly configured DNS forwarding proxies in NIOS, Infoblox recommends that you keep this option selected until you have verified that the NIOS proxies are functioning properly. In the Cloud Services Portal, go to **Manage -> On-Prem Hosts** to ensure that the statuses for the NIOS proxies that you have registered are active.



#### Note

- If you have upgraded to NIOS 8.5.x with DNS forwarding proxy enabled on any node, Infoblox recommends that you do not remove the on-prem hosts from the Cloud Services Portal. This is because NIOS preserves the access key during the upgrade and the NIOS Grid member connects to the Cloud Services Portal using the same access key.
- You must create a join token to authenticate a virtual DNS forwarding proxy for establishing a connection to the cloud. For more information on creating a join token, see the [Managing Join Tokens for On-Prem Hosts](#) section in the BloxOne Threat Defense documentation.
- If you have upgraded NIOS, the value of the **Access Key** field is the same as the API key that is displayed in the Cloud Services Portal.

## Enabling DNS Forwarding Proxy to Work with DNSSEC

DNS forwarding proxy does not work with DNSSEC in case a request was redirected by BloxOne Threat Defense. To enable DNS forwarding proxy to work with DNSSEC, perform the following steps:

1. Enable DNS Forwarding Proxy on NIOS by clicking **Manage -> On-Prem Hosts** in the Cloud Services Portal.
2. On Grid Manager, Disable the **Fall back to the default resolution process if BloxOne Threat Defense Cloud does not respond** option.
3. Enable DNSSEC validation as described in [Enabling DNSSEC Validation](#).
4. Remove trust anchors if any. To configure trust anchors, see [Enabling DNSSEC Validation](#).

## Controlling DNS Queries

By default, the NIOS appliance responds to DNS queries from any IP address. You can create a list of queriers to which the appliance is allowed to respond; restricting it to specific networks, IP addresses, and remote servers that present specified TSIG (transaction signature) keys. When using TSIG keys, it is important that the appliances and servers involved with the authentication procedure use NTP (Network Time Protocol) for their time settings (see [Using NTP for Time Settings](#)).

In addition, you can also configure the appliance to respond to recursive queries. A recursive query requires the appliance to return requested DNS data, or locate the data through queries to other servers. Recursion is disabled by default. If you enable this feature, you can also create a list of allowed recursive queries. For information about allowing recursion, refer to [Enabling Recursive Queries](#).

You can create a list of allowed queriers for the Grid and for individual Grid members.

## Specifying Queries

To configure a list of allowed queriers for the Grid or for a member:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click **Toggle Advanced Mode**, select the **Queries** tab.
3. In the Allow queries from section, select one of the following:
  - **Any:** Select this if you do not want to configure access control for DNS queries. The appliance allows queries from all clients. This is selected by default.
  - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this, the appliance allows clients that have the **Allow** permission to send and receive DNS queries. You can click **Clear** to remove the selected named ACL.
  - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows:
    - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address of the remote querier. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
    - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
      - **Permission:** Select **Allow** or **Deny** from the drop-down list.
    - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
      - **Permission:** Select **Allow** or **Deny** from the drop-down list.
    - **TSIG Key:** In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
      - **Key name:** Enter a meaningful name for the key, such as a zone name or the name of the remote name server. This name must match the name of the same TSIG key on other name servers.
      - **Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.
      - **Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.
    - **Any Address/Network:** Select to allow or deny queries from any IP addresses.

After you have added access control entries, you can do the following:

  - Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
  - Reorder the list of ACEs using the up and down arrows next to the table.
  - Select an ACE and click the Edit icon to modify the entry.
  - Select an ACE or multiple ACEs and click the Delete icon to delete the entries.
4. Save the configuration.

## Enabling Recursive Queries

You can enable the appliance to respond to recursive queries and create a list of allowed networks, IP addresses, and remote servers that present specified TSIG (transaction signature) keys. When using TSIG keys, it is important that the appliances and servers involved with the authentication procedure use NTP (Network Time Protocol) for their time settings (see [Using NTP for Time Settings](#)).

A recursive query requires the appliance to return requested DNS data, or locate the data through queries to other



servers. When a NIOS appliance receives a query for DNS data it does not have and you have enabled recursive queries, it first sends a query to any specified forwarders. If a forwarder does not respond (and you have disabled the **Use Forwarders Only** option in the **Forwarders** tab of the *Member DNS Properties* editor), the appliance sends a non-recursive query to specified internal root servers. If no internal root servers are configured, the appliance sends a non-recursive query to the Internet root servers. For information on specifying root name servers, see [About Root Name Servers](#).

You can enable recursion for a Grid, individual Grid members, and DNS views. For information about enabling recursion in a DNS view, see [Configuring DNS Views](#). If you do not enable recursion, the appliance denies recursive queries from all clients.

## Enabling Recursion

To enable recursion and create a list of recursive queriers:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox -> Edit icon. To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click **Toggle Advanced Mode**, select the **Queries** tab.
3. Click **Allow recursion**, and then in the *Allow recursive queries from* section, select one of the following:
  - **Any:** Select this if you want to configure access control for recursive queries. When you select **Any**, the appliance allows recursive queries from all clients. This is selected by default.
  - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this, the appliance allows clients that have the **Allow** permission to send and receive recursive DNS queries. You can click **Clear** to remove the selected named ACL.
  - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
    - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address of the remote querier. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
    - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
      - **Permission:** Select **Allow** or **Deny** from the drop-down list.
    - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
      - **Permission:** Select **Allow** or **Deny** from the drop-down list.
    - **TSIG Key:** In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
      - **Key name:** Enter a meaningful name for the key, such as a zone name or the name of the remote name server. This name must match the name of the same TSIG key on other name servers.
      - **Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.
      - **Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.
    - **Any Address/Network:** Select to allow or deny queries from any IP addresses.

After you have added access control entries, you can do the following:

  - Select the ACEs that you want to consolidate and put into a new named ACL. Click the **Create new named ACL** icon and enter a name in the *Convert to Named ACL* dialog box.

The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.

- Reorder the list of ACEs using the up and down arrows next to the table.
- Select an ACE and click the Edit icon to modify the entry.
- Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.

4. Save the configuration.

## Configuring Resolver Queries Timeout

You can configure the amount of time that a recursive query will wait for a response before timing out. The default timeout behavior is to wait for 30 seconds before timing out.

To configure the resolver queries timeout value:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click **Toggle Advanced Mode**, select the **Queries** tab and complete the following:
  - **Resolver queries timeout:** Specify the maximum time allowed for a recursive query to wait for a response before timing out. You can enter either 0 or a value between 10 and 30 seconds. Setting the timeout value to 0 returns to the default timeout behavior, which is to wait for 30 seconds before timing out.
3. Save the configuration.

## Restricting Recursive Client Queries

By default, the appliance can serve up to 1,000 outstanding recursive client queries. You can change this default value according to your business needs. After you configure the recursive client queries limit, you can enable the appliance to send SNMP traps for recursive queries. Enabling SNMP traps for recursive clients can help you identify possible flood attacks on the DNS recursive server. The appliance sends SNMP traps when the number of recursive client queries exceeds the configured thresholds. For information about how to set the threshold and reset values, see [Defining Thresholds for Traps](#).

1. From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.
2. In the *Member DNS Properties* editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Advanced** subtab of the **Queries** tab.
4. Select the **Limit number of recursive clients** to option and enter a number. By default, the appliance is allowed to serve up to 1000 concurrent clients that send recursive queries. You can change this default according to your business needs from between 0 to 40000.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Enabling Recursive Resolution Using EDNS Client Subnet (ECS) Option

The EDNS Client Subnet (ECS) option is a DNS extension you use to optimize recursive resolution for query sources that are not topologically close to the recursive resolvers. When you enable ECS for recursive resolution, the appliance includes subnet information of the host that originates a DNS query. Thus, your recursive resolver can perform geotargeting by passing the subnet information to authoritative servers so that the response will be more optimized for the end clients. For example, when you enable ECS and/or ECS forwarding on your recursive resolver, CDNs (Content Delivery Networks) can deliver content faster and more efficiently to the end user by providing information about the end user's subnet to the authoritative DNS server operated by the CDNs.

You can enable the NIOS appliance to handle recursive queries using ECS option and enable ECS forwarding support at the Grid level. You can then add whitelisted zone names that are subject to ECS recursion and specify the source prefix length for IPv4 and IPv6 addresses. Make sure you enter only apex zones. Example: foo.com, corpxyz.com, etc. The whitelisted zone name indicates the zone to which ECS tagged queries must be sent.

Note the following while adding whitelisted zone names:

- ECS options are sent only when the name being queried and the apex of the zone being queried both match ECS zones. For example, if the zone "foo.com" contains a subdomain "www.foo.com", then you must configure



"foo.com" as ECS zone and not "www.foo.com". The latter configuration might result in no ECS queries being sent, because the apex zone, "foo.com" does not match with "www.foo.com".

- Queries for subdomains of the specified zone name, with prefix lengths greater than the specified prefix length is not applicable for the subdomains of the specified zone name. For example, if you specify "foo.com" with IPv4 prefix length 20, then IPv4 queries with prefix length greater than 20 is not applicable for the subdomains of "foo.com".
- You can exclude certain subdomains by adding a leading exclamation mark (!) to the subdomain name. Example: ! foo.example.org, ! test.foo.com, etc.

## Guidelines for Using ECS and ECS Forwarding

The following are the guidelines for using ECS and ECS forwarding:

- When recursive ECS is enabled, the appliance applies ECS handling for queries that meet both of the following criteria:
  - If the source prefix length is not set to zero.
  - If the query name matches one of the listed whitelisted zone names.
- If you enable ECS forwarding, all queries that contain a valid ECS option will be forwarded to the authoritative server.
- Queries with the source prefix length set to zero will be forwarded unchanged, regardless of whether ECS forwarding is enabled or disabled.
- When recursive ECS and ECS forwarding are enabled, then response to queries that contain a valid ECS option with a non-zero source prefix length will contain an ECS option.
- When recursive ECS is enabled and ECS forwarding is disabled, and if the original query contains a valid ECS option with a non-zero source prefix length, then the resolver returns a REFUSED response.

To enable recursive ECS and configure DNS resolver parameters, complete the following:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the *Grid DNS Properties* editor, click the **Advanced** subtab of the **Queries** tab and complete the following:
  - **Enable Recursive ECS**: Select this checkbox to enable recursive resolution using ECS. This is disabled by default. If recursive ECS is enabled, the appliance applies ECS handling for queries that meet both of the following criteria:
    - If the source prefix length is not set to zero.
    - If the query zone name is listed in the whitelisted domains.
  - **Enable ECS Forwarding**: Select this checkbox to enable ECS forwarding. If you enable ECS forwarding, all queries containing a valid ECS option will be forwarded to the authoritative server.  
**Note**: Queries with the source prefix length set to zero will be forwarded unchanged, regardless of whether ECS forwarding is enabled or disabled.
  - **Query Zone Permissions**: Click the Add icon to add a list of query zone names that are subject to ECS recursion and the corresponding permission. Grid Manager adds a row to the table. Complete the following:
    - **Zone Name**: Enter the zone name.
    - **Permission**: Select **Allow** or **Deny** from the drop-down list.
  - **IPv4SourcePrefix**: Specify the IPv4 source prefix length. You can enter a value between 1 and 24. The default value is 24.
  - **IPv6SourcePrefix**: Specify the IPv6 source prefix length. You can enter a value between 1 and 56. The default value is 56.

## Enabling DNS Fault Tolerant Caching

When an authoritative DNS server experiences an outage, all web sites served by the DNS server become inaccessible. Enabling the DNS fault tolerant caching option allows users to access the web sites served by the DNS server despite the DNS server outage. When you enable the DNS fault tolerant caching option, DNS records are retained in the recursive cache even after they expire. Whenever recursive query times out or returns a SERVFAIL response, the appliance returns the cached response to the client instead of the SERVFAIL response.

When you enable DNS fault tolerant cache, you can also specify the TTL (time-to-live) and timeout settings for the expired records. TTL specifies the time duration for which the expired record is retained in the recursive cache. Setting a high TTL might cause the client to use incorrect data for a longer duration. Conversely, setting a low TTL renders more

current cached data, but also increases the traffic on your network. The expired record is deleted from the recursive cache after the specified timeout duration.

Only DNS members with recursion enabled can support this feature. You can enable this feature at the Grid level and override it at member level with recursion enabled. For information on enabling recursion for a Grid or member, see [Enabling Recursive Queries](#). Note that DNS fault tolerant caching does not work when you set the **DNS Resolver Type** to **Unbound**.

To enable DNS fault tolerant caching, complete the following:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
To override Grid settings, click **Override** next to it and complete the appropriate fields.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click **Advanced** subtab of the **Queries** tab and complete the following:
  - **Enable Fault Tolerant Caching:** Select this checkbox to enable the retention of expired records in the recursion. When you enable this option, the appliance retains the expired records in the recursive cache. Whenever recursive queries times out or returns a SERVFAIL response, the appliance returns the cached response to the client instead of the SERVFAIL response. This is disabled by default.
  - **Expired Record TTL:** Specify the time duration that the appliance must serve the expired records from the recursive cache before attempting to refresh the records. The default is five seconds. Select the time period in minutes, hours, or days from the drop-down list.
  - **Expired Record Timeout:** Specify the time duration that the appliance waits before deleting the expired records from recursive cache. The default is 24 hours. Select the time period in minutes, hours, or days from the drop-down list.
3. Save the configuration.

## Controlling AAAA Records for IPv4 Clients

By default, the NIOS appliance returns resource records, including AAAA records, in response to DNS queries. You can enable the appliance to filter and remove AAAA records in response to queries received over IPv4 for each name server and DNS view. This feature is useful in a configuration where a client issues a DNS query over IPv4 when it does not have the ability to use an IPv6 address. When a response returns an IPv6 address however, the client that sends the query over an IPv4 transport would lose connectivity. By enabling AAAA filtering, you can configure your name server not to return AAAA records to clients that request queries over an IPv4 transport. Presumably, these clients then re-query the name server for A records for the same domain name.

Depending on your configuration, the appliance can remove AAAA records for all queries over IPv4 (even when DNSSEC is enabled), or only for queries that are not DNSSEC-signed. You can also create a list of IPv4 networks and addresses to which the appliance applies AAAA filtering and vice versa. You can enable and configure AAAA filtering for the Grid, members, and DNS views.

To control whether you want the appliance to return AAAA records for queries sent over IPv4, you must first enable AAAA filtering, and then create a list of IPv4 networks and addresses that allow or deny AAAA filtering from the appliance, as described in [Enabling AAAA Filtering](#).



### Note

An AAAA record is filtered only when there is also an A record for the same domain name. In this case, the appliance still sends a response, but without any AAAA or A record in it. When a client queries for an AAAA record and there is no corresponding A record for it, the appliance returns the AAAA record even if you have enabled AAAA filtering for this client.

## Enabling AAAA Filtering

To enable AAAA filtering and configure a list of IPv4 networks and addresses:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab > **Members** tab -> member checkbox -> Edit icon.  
**DNS View:** From the **Data Management** tab, click the **DNS** tab -> **Zones** tab > *dns\_view* checkbox -> Edit icon.  
 To override Grid settings, click **Override** and complete the appropriate fields.
2. In the editor, click the **Queries** tab and complete the following:
 

**Enable AAAA Filtering:** From the drop-down list, select one of the following:

  - **Break DNSSEC:** Select this to remove AAAA records in response to queries sent over IPv4, including those that are signed by DNSSEC.  
 Note: Be aware that when you select this option, DNSSEC configuration will no longer be in effect.
  - **No:** Select this to disable AAAA filtering for queries over IPv4. When you select this, the appliance returns AAAA records in response to all DNS queries issued over IPv4. This is selected by default.
  - **Yes:** Select this to enable AAAA filtering for queries over IPv4. When you select this, the appliance removes AAAA records in response to all DNS queries issued over IPv4, except for DNSSEC-signed requests.
3. In the *AAAA Filtering* section, select one of the following:
  - **None:** Select this if you want to configure access control for AAAA filtering. The appliance allows all clients to issue DNS queries over IPv4 when they do not have the ability to use IPv6 addresses. This is selected by default.
  - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this, the appliance allows clients that have the **Allow** permission can filter AAAA responses. You can click **Clear** to remove the selected named ACL.
  - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
    - **IPv4 Address:** Select this to add an IPv4 address. Click the **Value** field and enter the IP address of the client. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list. When you select **Allow**, the appliance applies AAAA filtering and removes AAAA records in response to queries sent by the specified IPv4 address. When you select **Deny**, the appliance does not apply AAAA filtering and thus returns AAAA records.
    - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
      - **Permission:** Select **Allow** or **Deny** from the drop-down list.
  - **Any Address/Network:** Select to allow or deny AAAA filtering from any IP addresses.  
 After you have added access control entries, you can do the following:
    - Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
    - Reorder the list of ACEs using the up and down arrows next to the table.
    - Select an ACE and click the Edit icon to modify the entry.
    - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.



#### Note

If you do not enter any addresses or networks in the table, the appliance applies AAAA filtering to all IPv4 clients. In other words, the appliance removes AAAA records in responses to all queries sent over IPv4.

## About NXDOMAIN Redirection

When a DNS member with recursion enabled receives a recursive query for data for which it is not authoritative, it locates the data through queries to other servers. If the query is for a non-existent domain name, the DNS member receives an NXDOMAIN response from the authoritative name server, which the member then forwards to the DNS client. An NXDOMAIN response contains a "Name Error" RCODE, signifying that the domain name referenced in the query does not exist. (For information, you can refer to *RFC 1035, Domain Names — Implementation and Specification*.) You can install a Query Redirection license on a recursive DNS member to control its response to queries for A records of non-existent domain names and other domain names that you specify. After the license is installed, Grid Manager displays the **NXDOMAIN Rulesets** tab where you can create rules that specify how a DNS member responds to queries for A/AAAA records for certain domain names and non-existent domain names. Each rule contains a domain name specification, and the action of the DNS member when the domain name in the query matches that in the rule. After you create the rules, you then enable the NXDOMAIN redirection feature and list the IP addresses that are included in the synthesized responses.

Recursive DNS members can redirect responses to queries for A/AAAA records only. DNS members resolve queries for all other records as they normally would.

In addition, you can enable DNS members to log queries that match rules with an action of "Redirect" or "Modify". You can view the logs in the *Syslog* viewer. The logs include the queried domain name, source IP address, the pattern of the matched rule, and the name of the corresponding ruleset.

When DNSSEC is enabled on the Infoblox DNS server, it does not redirect DNS clients that request DNSSEC data for a non-existent domain name. Instead, it returns an authenticated negative response in the form of an NSEC or NSEC3 RR. (For information about DNSSEC, see [Configuring DNSSEC](#).) If DNSSEC is not enabled, the appliance ignores the request for DNSSEC data and redirects the clients.

To apply the configured NXDOMAIN rules regardless of whether a DNS query requests DNSSEC data, configure the appliance accordingly. For more information about how to configure this, see [Applying Policies and Rules to DNS Queries that Request DNSSEC Data](#).

You can enable NXDOMAIN redirection at the Grid, member, and DNS view levels. Only recursive DNS servers can redirect DNS clients. Non-recursive DNS members do not redirect DNS clients. For information on enabling recursion on a DNS member, see [Enabling Recursive Queries](#).

Note that if both NXDOMAIN redirection and the blacklisting feature are enabled, the DNS member applies the blacklist rulesets before the NXDOMAIN rulesets. For information about blacklisting domain names, see [Blacklists](#).

## About NXDOMAIN Rulesets

An NXDOMAIN ruleset is a list of rules that a DNS member uses to determine its response to recursive queries for A records it does not have. Each rule consists of a domain name specification or pattern, and an associated action. Domain names can contain any printable character. You can use certain metacharacters to create domain name patterns that are used to match the domain names in DNS queries. Pattern matching is case-insensitive. Patterns support the following metacharacters:

- Use the caret character (^) to indicate the beginning of a pattern. For example, **^foo** matches **foo.com** but not **barfoo.com**. The caret character has a special meaning only if it is specified at the beginning of a pattern.
- Use the dollar sign character (\$) to indicate the end of a pattern. The dollar sign character has a special meaning, only if it is specified at the end of the pattern. For example, **.com\$** matches **corpxyz.com** but not **corpxyz.com.net**. When the pattern contains a \$ at the end, NIOS automatically adds a period (.) before the \$. For example, if you enter **.com\$**, NIOS saves it as **.com.\$**. The period indicates that the pattern specifies a complete domain name that ends with the root label.
- Use the asterisk character (\*) as a wildcard that can match zero or more characters in one or more labels of a domain name. For example, **x\*oy** matches **xfooy.com**, but not **xfoobary**. A pattern that contains a single asterisk (\*) (or an equivalent expression, such as **\*\*A\${\*}**) matches any domain name.
- Use the backslash character (\) with one of the metacharacters (\$, ^, \*\* and \*) to remove their special meaning. If \ is followed by any other character, that character is taken as an ordinary character, as if \ is not present. For example, **foo\\.bar** matches **foo.bar**, and **\\*** matches a literal asterisk in a domain name.

No other characters have any special meaning. Note in particular that the period character (".") only matches a period used as a separator in a domain name.

The action specifies how the DNS member responds when a domain name in a query matches a pattern. The action can be one of the following: **Pass**, **Modify** or **Redirect**.

- **Pass:** The DNS member resolves the query and forwards the response to the DNS client, even if it is an NXDOMAIN response.
- **Modify:** The DNS member resolves the query and forwards the response to the DNS client, only if it is not an NXDOMAIN response. But if the member receives an NXDOMAIN response, it sends the client a synthesized response that includes predefined IP addresses.
- **Redirect:** The DNS member does not resolve the query. Instead, it sends the client a synthesized response that includes predefined IP addresses.

You can configure multiple rulesets. The DNS member applies the rulesets and their rules in the order in which they're specified in the configuration. If multiple rulesets contain rules with duplicate patterns, the DNS member applies the first rule it encounters and ignores the other rules.

## Examples

The following example illustrates how the appliance applies NXDOMAIN rulesets. Ruleset 1:

Pattern	Action
a1.corpxyz.com	PASS
*.corpxyz.com	REDIRECT

- If the DNS member receives a query for a1.corpxyz.com, it resolves the query and forwards the response, even if it is an NXDOMAIN response, to the client. Note that if the order of the rules was switched, the DNS client would have been redirected immediately, because the domain name a1.corpxyz.com matches the \*.corpxyz.com pattern.
- If the DNS member receives a query for b1.corpxyz.com, the member immediately redirects the DNS client to the specified IP address because the domain name in the query matches the second rule.
- If the DNS member receives a query for b1.corp200.com, it resolves the query because the domain name does not match any rule. If the DNS member receives an A record from an authoritative server, the member forwards the response to the client. However, if the member receives an NXDOMAIN response, it redirects the DNS client to the specified IP address.

In the following example, the rules redirect queries for dotted domain names that do not have ".com" As shown in the example, an explicit PASS rule is required at the end.

Ruleset 2:

Pattern	Action
*.com	PASS
.*.\$	MODIFY
*	PASS

- If the DNS member receives a query for corpxyz.com which matches the pattern "\*.com", the member resolves the query and forwards the response, even if it is an NXDOMAIN response, to the client.

- If the DNS member receives a query for corpxyz.org, which matches the pattern ".\*\$", the member resolves the query. If the member receives an NXDOMAIN response, it redirects the client to the specified IP address. If the member receives a non-NXDOMAIN response, it forwards the response to the client.
- If the DNS member receives a query for corp200, the member resolves the query and forwards the response to the client.

## NXDOMAIN Redirection Guidelines

The following summarizes how a DNS member responds to a query for an A record when the NXDOMAIN feature is enabled:

- If there are no rulesets configured, the DNS member queries other name servers.
  - If the DNS member receives a non-NXDOMAIN response from an authoritative server, it forwards the response to the DNS client.
  - If the DNS member receives an NXDOMAIN response from an authoritative server, it redirect the DNS client.
- If rulesets are configured, the DNS member tries to match the domain name in the query with a domain name in the rules.
  - If the DNS member finds a match, it perform the action specified in the rule.
    - If the action is "Redirect", the DNS member redirect the DNS client.
    - If the action is "Pass", the DNS member queries other name servers and forwards the response to the DNS client.
    - If the action is "Modify", the DNS member queries other name servers. If it receives a non-NXDOMAIN response, it forwards the response to the DNS client; if it receives and NXDOMAIN response, it redirects the DNS client.
  - If the DNS member does not find a match, the DNS member queries other name servers.
    - If the DNS member receives a non-NXDOMAIN response, it forwards the response to the DNS client.
    - If the DNS member receives an NXDOMAIN response from an authoritative server, it redirects the DNS client.

Note that if an A record with a dotted hostname is added to an authoritative zone through a dynamic DNS update, and that A record should actually belong in an existing delegation, the appliance may not redirect a query for that A record according to the Blacklist and NXDOMAIN guidelines.

## Configuring NXDOMAIN Redirection

To enable NXDOMAIN redirection and configure its properties:

1. Configure NXDOMAIN rulesets. You can create NXDOMAIN rulesets through Grid Manager, as described in [Creating Rulesets](#) below. You can also specify the rulesets in a CSV file and import the file to the Grid, as described in [Importing and Exporting Data using CSV Import](#).
2. Enable this feature and specify the redirection IP addresses, as described in [Enabling NXDOMAIN Redirection](#) below.

## Creating Rulesets

To create a ruleset:

1. From the **Data Management** tab -> **DNS** tab -> **NXDOMAIN Rulesets** tab, click the Add icon.
2. In the **NXDOMAIN Ruleset** wizard, complete the following and click **Next**:
  - **Name**: Enter a name for the ruleset.
  - **Comment**: You can enter additional information.
  - **Disable**: You can disable this ruleset for use later on. The appliance ignores disabled rulesets.
3. Click the Add icon to add a rule to the ruleset table.
  - In the **Pattern** column, enter a domain name or pattern, using the guidelines specified in [About NXDOMAIN Rulesets](#).
  - In the **Action** column, select **PASS**, **REDIRECT** or **MODIFY**.

- In the **Order** column, NIOS automatically displays the number of the entry in the list. The appliance applies the rules in the order they are listed. You can order the list as follows:
    - Use the up and down arrows to move rules up or down on the list.
    - Use the go-to-top or go-to-bottom arrow to move a rule to the top or bottom of the list.
    - Change the Order number of a rule to move it to the desired location.
    - Delete a rule by selecting it and clicking the Delete icon.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Managing NXDOMAIN Rulesets

To view NXDOMAIN rulesets, navigate to the **Data Management** tab -> **DNS** tab -> **NXDOMAIN Rulesets** tab. The panel lists the configured rulesets and their associated comments. You can also display the Disabled column which indicates which rulesets are disabled. From this panel, you can do the following:

- Add more rulesets, as described in the preceding section, [Creating Rulesets](#).
- Edit a ruleset, by clicking its checkbox and clicking the Edit icon. You can set the following in the NXDOMAIN Ruleset editor:
  - In the **General Basic** tab, you can change entries in any of the fields.
  - In the **Rules** tab, you can do the following:
    - Add a rule by clicking the add icon and specifying the pattern and action.
    - Change the pattern or action of a rule, by clicking in the appropriate row.
    - Delete a rule by clicking its checkbox and clicking the Delete icon.
    - Move rules up and down, by using the arrows.
    - In the **Permissions** tab, you can set admin permissions for the ruleset. For information about admin permissions, see [Managing Administrators](#).
  - Delete a ruleset, by clicking its checkbox and clicking the Delete icon.

## Enabling NXDOMAIN Redirection

Only DNS members with recursion enabled can support NXDOMAIN redirection.

You can enable this feature at the Grid level, and override it for a member or DNS view with recursion enabled. You must specify at least one IP address as the redirection destination. You can specify different redirection IP addresses and rulesets for each Grid member or DNS view, and you can also define members that do not provide redirection. This is useful when you want to define a set of "opt out" servers for DNS clients that do not want to be redirected.

You can also enable the DNS member to log queries that match rules with an action of "Redirect" or "Modify". The logs include the queried domain name, source IP address, the pattern of the matched rule, and the name of the corresponding ruleset. The DNS member does not log queries that matched rules with an action of "Pass".

To enable NXDOMAIN redirection:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
**DNS View:** From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns\_view* checkbox -> Edit icon.  
**Standalone DNS:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **System DNS Properties**.  
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. If the *Grid DNS Properties* or *Member DNS Properties* editor is in basic mode, click **Toggle Advanced Mode**.
3. Click **NXDOMAIN** and complete the following:
  - **Enable NXDOMAIN redirection (recursive members only):** Select this option to enable recursive DNS members to synthesize their responses to DNS queries for A records.
  - **Rulesets:** Click the Add icon to add an NXDOMAIN ruleset. Use the up and down arrows to move rulesets up and down in the list. The appliance applies them in the order they are listed.
  - **Redirect to IPv4 addresses:** Click the Add icon and enter the IPv4 addresses that the DNS server includes in its synthesized response for A type queries.
  - **Redirect to IPv6 addresses:** Click the Add icon and enter the IPv6 addresses that the DNS server includes in its synthesized response for AAAA type queries.  
 Note that you can add up to 12 IP addresses, combination of both IPv4 and IPv6, for NXDOMAIN redirection.



- **TTL:** Specify how long the DNS client caches the A record with the redirected IP address.
  - **Log redirected queries:** Select this checkbox to log the redirected queries to syslog.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring DNS over TLS and DNS over HTTPS Services

DNS queries and responses sent over port 53 without encryption are vulnerable to spoofing and eavesdropping. This issue is addressed in NIOS appliances that have DNS over TLS (Transport Layer Security) and DNS over HTTPS services enabled. These features encrypt DNS queries and responses to secure communication between a DNS server and a DNS client.

This topic details the requirements that NIOS appliances must meet for enabling the DNS over TLS and DNS over HTTPS services and has instructions to configure these services. The sections covered in this topic are as follows:

- [Licensing and Certificate Requirements](#)
- [Base Configuration Requirements](#)
  - [Configuration Requirements if Parental Control is Enabled](#)
- [Supported Cipher Suites](#)
- [Limitations and Recommendations for DNS over TLS and DNS over HTTPS](#)
- [DNS over TLS](#)
  - [Configuring DNS over TLS](#)
  - [CLI Support for DNS over TLS](#)
- [DNS over HTTPS](#)
  - [Configuring DNS over HTTPS](#)
  - [CLI Support for DNS over HTTPS](#)

### Licensing and Certificate Requirements

DNS over TLS and DNS over HTTPS require the vDCA (virtual DNS Cache Acceleration) or vADP (virtual Advanced DNS Protection Software) service to be licensed and enabled. If the DNS Cache Acceleration and/or Advanced DNS Protection Software services are not enabled, the DNS over TLS and DNS over HTTPS features will not work even if they are enabled. For more information about DNS Cache Acceleration and Advanced DNS Protection Software (threat protection), see [Configuring DNS Cache Acceleration](#) and [About Infoblox Advanced DNS Protection](#) respectively.

The DNS over TLS or the DNS over HTTPS service uses the same self-signed certificate that NIOS generates for HTTPS communication when it first starts. You can also generate a certificate signing request (CSR) and use it to obtain a signed certificate from your own trusted certificate authority (CA). For more information, see [Generating Certificate Signing Requests](#).

The certificate is provisioned for each member. For more information about certificates, see [Managing Certificates](#).



#### Note

NIOS generates a new self-signed certificate when the host name or the IP address of the member is changed or when a Grid Master Candidate is promoted. If the DNS over TLS or DNS over HTTPS feature is enabled on a member, then every time a new self-signed certificate, HTTPS certificate, or a CA certificate is generated, the DNS over TLS service or the DNS over HTTPS service (depending on which feature is enabled) automatically restarts to upload the new certificate.

### Base Configuration Requirements

NIOS appliances must have the required base memory configuration to enable the DNS over TLS and the DNS over HTTPS features on their members. If the appliances do not meet the required criteria, the options to configure these



features are not displayed in the *Member DNS Properties* editor. The following table lists the base configuration required for enabling these features on IB-FLEX appliances.



**Warning**

The numbers in the following tables are for IB-FLEX appliances only. For information about CPU and memory requirements of NIOS appliances other than IB-FLEX, see the *NIOS Release Notes*.

IB-FLEX Flavor Configuration	Total CPU	Total System Memory in GB (With virtual Advanced DNS Protection Software only)	Total System Memory in GB (With virtual DNS Cache Acceleration and virtual Advanced DNS Protection Software)	Maximum Number of Concurrent Sessions Supported	Grid Master Capable
Small recursive DNS (with acceleration)	10	32	32	For vDCA only: 120,000 For vADP only: 50,000 For vDCA and vADP: 120,000	No
Medium recursive DNS (with acceleration)	16	64	40	For vDCA only: 150,000 For vADP only: 60,000 For vDCA and vADP: 150,000	No
Large recursive DNS (with acceleration)	26	80	50	For vDCA only: 240,000 For vADP only: 80,000 For vDCA and vADP: 240,000	No

The following table lists the maximum number of concurrent sessions supported by different NIOS appliance models (physical and virtual). For information about CPU and memory requirements, see the NIOS Release Notes.



**Note**

If the available memory does not meet the requirement defined in the above table, you may observe unexpected behavior. Infoblox recommends that you allocate slightly more memory to ensure that memory associated with the hypervisor is also accounted for.

NIOS Appliance (Physical and Virtual)	Maximum Number of Concurrent Sessions Supported
IB-14x5	For vADP only: 50,000

NIOS Appliance (Physical and Virtual)	Maximum Number of Concurrent Sessions Supported
IB-22x5	For vDCA only: 150,000 For vADP only: 60,000 For vDCA and vADP: 150,000
IB-40x5	For vDCA only: 240,000 For vADP only: 80,000 For vDCA and vADP: 240,000



#### Note

In an HA setup, ensure that both the active and passive nodes have the memory configuration required to enable the DNS over TLS or the DNS over HTTPS feature. If you enable the feature on an active node that has the required memory footprint but the passive node does not, then in case of a failover, the DNS over TLS or the DNS over HTTPS service does not start on the new active node. Therefore, requests coming to the DNS over TLS or the DNS over HTTPS stream are not honored.

#### Configuration Requirements if Parental Control is Enabled

NIOS appliances require additional memory if you intend to run DNS over TLS and/or DNS over HTTPS along with the Parental Control features such as proxy RPZ passthru, DCA subscriber query count logging, and DCA subscriber allowed and blocked listing simultaneously. The following table lists the base configuration required on IB-FLEX appliances for configuring these features simultaneously:

IB-FLEX Flavor Configuration	Total CPU	Total System Memory in GB (With virtual DNS Cache Acceleration only)	Total System Memory in GB (With virtual DNS Cache Acceleration and virtual Advanced DNS Protection Software)	Maximum Number of Concurrent Sessions Supported	Grid Master Capable
Medium recursive DNS (with acceleration)	16	64	64	For vDCA only: 150,000 For vADP only: 60,000 For vDCA and vADP: 150,000	No
Medium-Large recursive DNS (with acceleration)	16	86	86	For vDCA only: 150,000 For vADP only: 60,000 For vDCA and vADP: 150,000	No

IB-FLEX Flavor Configuration	Total CPU	Total System Memory in GB (With virtual DNS Cache Acceleration only)	Total System Memory in GB (With virtual DNS Cache Acceleration and virtual Advanced DNS Protection Software)	Maximum Number of Concurrent Sessions Supported	Grid Master Capable
Large recursive DNS (with acceleration)	26	100	100	For vDCA only: 240,000 For vADP only: 80,000 For vDCA and vADP: 240,000	No



**Note**

- When a NIOS appliance does not have the required base memory configuration, if you try to enable and run DNS over TLS, DNS over HTTPS, and Parental Control features simultaneously, all of these features will be disabled.
- For information about CPU and memory requirements of NIOS appliance models other than IB-FLEX, see the NIOS Release Notes.

### Supported Cipher Suites

From NIOS 8.6.1 onwards, the DNS over TLS and DNS over HTTPS features support only TLS version 1.2 and TLS version 1.3 cipher suites. The cipher suite order preference is configured to improve the throughput in DNS over TLS and DNS over HTTPS communication.

Cipher suites supported for TLS 1.2 are as follows:

- ECDHE-ECDSA-CHACHA20-POLY1305
- ECDHE-RSA-CHACHA20-POLY1305
- ECDHE-ECDSA-AES128-GCM-SHA256
- ECDHE-ECDSA-AES256-GCM-SHA384
- ECDHE-ECDSA-AES128-SHA256
- ECDHE-ECDSA-AES256-SHA384
- ECDHE-RSA-AES128-GCM-SHA256
- ECDHE-RSA-AES256-GCM-SHA384
- ECDHE-RSA-AES128-SHA256
- ECDHE-RSA-AES256-SHA384
- DHE-RSA-AES128-GCM-SHA256
- DHE-RSA-AES256-GCM-SHA384
- DHE-RSA-AES128-SHA
- DHE-RSA-AES256-SHA
- DHE-RSA-AES128-SHA256
- DHE-RSA-AES256-SHA384

Cipher suites supported for TLS 1.3 are as follows:

- TLS\_CHACHA20\_POLY1305\_SHA256
- TLS\_AES\_128\_GCM\_SHA256
- TLS\_AES\_128\_CCM\_SHA256
- TLS\_AES\_128\_CCM\_8\_SHA256
- TLS\_AES\_256\_GCM\_SHA384

## Limitations and Recommendations for DNS over TLS and DNS over HTTPS

Consider the following limitations and recommendations when you enable the DNS over TLS and/or the DNS over HTTPS features:

- If an appliance configured with DNS over TLS or DNS over HTTPS has both vDCA and vADP running, the configuration is set to the DCA-first mode.
- TSIG queries for which responses are larger than the max EDNS/UDP buffer size are not supported.
- DNS queries coming with EDNS padding over port 53 are dropped.
- DNS over TLS and DNS over HTTPS features are not supported on unbound-based DNS servers.
- When DNS over TLS or DNS over HTTPS is enabled, queries decrypted at DNS over TLS or DNS over HTTPS that do not receive a response from the vDCA cache are forwarded to the recursive DNS engine over UDP. Therefore, rules added for TCP requests over TLS or HTTPS may not be honored. Infoblox recommends that you add the corresponding UDP-specific rules instead of only the TCP request rules.
- For NIOS 8.5.2 only: Infoblox recommends that you manually set the maximum packet size of both the UDP buffer and the EDNS buffer to 4096 bytes. If the packet size exceeds 4096, packets are dropped by the DNS over TLS or the DNS over HTTPS server. For more information about setting buffer sizes, see [Configuring the EDNS0 Buffer Size and UDP Buffer Size](#).
- DNS over TLS only:
  - The TLS versions that are currently supported by NIOS are TLS 1.2 and TLS 1.3.
  - DNS over TLS supports queries and responses from both DNS and DNS Cache Acceleration services.
  - DNS over TLS is not supported for recursive queries when performing upstream lookups.
  - DNS zone transfer requests over DNS over TLS are not supported.
  - For DNS over TLS clients that use systemd-resolved service, the Subject Alternative Name (SAN) must point to the IP address of the DNS service. By default, the self-signed certificates issued to Infoblox members do not meet this requirement. Therefore, for Infoblox to support systemd-resolved, you must install certificates that include SAN IP address from a trusted certificate authority.
- DNS over HTTPS only:
  - DNS over HTTPS is supported on the HTTP/2 protocol.
  - DNS over HTTPS is supported only if the NIOS appliance has an MGMT interface set up. The DNS over HTTPS module listens on port 443 for interfaces other than MGMT and any incoming UI request to the MGMT interface is bypassed directly to the host.
  - When DNS over HTTPS is enabled on a member, HTTP redirection from the member to its Grid Master is disabled.

## DNS over TLS

NIOS appliances that support DNS Cache Acceleration or Advanced DNS Protection Software, include the DNS over TLS capability that helps increase DNS security and privacy. When you enable the DNS over TLS feature, DNS traffic is encrypted through the TLS protocol to prevent eavesdropping and tampering of DNS data. This feature is supported on both recursive and authoritative DNS servers only through port 853. It is available only for Grid members and for standalone systems. It supports the processing of multiple DNS queries/responses over a single TLS session.

You can configure and run the DNS over TLS service on a member only when the following prerequisites are met:

- Either the accelerated DNS Cache Acceleration (vDCA) or the Advanced DNS Protection Software (vADP) service is enabled.
- The memory required to support the DNS over TLS feature is available. For more information, see Base Configuration Requirements below.

## Configuring DNS over TLS

To configure the DNS over TLS feature, complete the following steps:

1. **Grid member:** On the **Data Management** tab, click the **DNS** tab -> **Members** tab, select the *member* checkbox, and then click the Edit icon.  
**Standalone system:** On the **Data Management** tab, click the **DNS** tab, expand the Toolbar, and then click **System DNS Properties**.

2. In the *Member DNS Properties* editor/*System DNS Properties* editor, click **Toggle Advanced Mode** if the editor is in basic mode.
3. On the **Queries** tab -> **Advanced** tab, select the **Enable DoT Service** checkbox to enable the DNS over TLS feature.  
Note the options for DNS over TLS feature are displayed only if the appliance has the memory footprint that is required to support the feature and has the virtual DNS Cache Acceleration or Advanced DNS Protection Software license installed. For more information, see Base Configuration Requirements below.
4. In the **Maximum Session Timeout** field, specify the maximum time in seconds a session can remain idle before it times out and closes. The default value is 60 seconds.  
If your DNS forwarders are located at different geographical locations or if the network latency is high, you may observe session timeouts. If so, Infoblox recommends that you set the **Maximum Session Timeout** to more than 60 seconds. Increasing the session duration may impact concurrent open sessions.
5. Save the configuration.
6. As prompted, manually reboot the member to enable the DNS over TLS feature.



#### Note

The DNS over TLS feature will not take effect until you reboot the member or the standalone system and ensure that either the DNS Cache Acceleration or Advanced DNS Protection Software service is running after the reboot.

### CLI Support for DNS over TLS

You can view the status of the DNS over TLS service, configuration, and details of active sessions using the following commands:

- *show dns-over-tls-status*
- *show dns-over-tls-config*
- *show dns-over-tls-stats*

### DNS over HTTPS

NIOS appliances that support DNS Cache Acceleration or Advanced DNS Protection Software, include the DNS over HTTPS capability that helps increase DNS security and privacy. When you enable the DNS over HTTPS feature, DNS traffic is encrypted through the HTTPS protocol to prevent eavesdropping and tampering of DNS data. This feature is supported on both recursive and authoritative DNS servers only through port 443. It is available only for Grid members and standalone systems. The feature supports the processing of multiple DNS queries/responses over a single TCP session.

You can configure and run the DNS over HTTPS service on a NIOS appliance only when the following prerequisites are met:

- An MGMT interface is set up.
- The memory required to support the DNS over HTTPS feature is available. For more information, see Base Configuration Requirements below.
- Either the accelerated DNS Cache Acceleration (vDCA) or the Advanced DNS Protection Software (vADP) service is enabled.

### Configuring DNS over HTTPS

To configure the DNS over HTTPS feature, complete the following steps:

1. **Grid member:** On the **Data Management** tab, click the **DNS** tab -> **Members** tab, select the *member* checkbox, and then click the Edit icon.  
**Standalone system:** On the **Data Management** tab, click the **DNS** tab, expand the Toolbar, and then click **System DNS Properties**.
2. In the *Member DNS Properties* editor/*System DNS Properties* editor, click **Toggle Advanced Mode** if the editor is in basic mode.

3. On the **Queries** tab -> **Advanced** tab, select the **Enable DoH Service** checkbox to enable the DNS over HTTPS feature.  
Note the options for DNS over HTTPS feature are displayed only if the appliance has the memory footprint that is required to support the feature and has the virtual DNS Cache Acceleration or Advanced DNS Protection Software license installed. For more information, see Base Configuration Requirements below.
4. In the **Maximum Session Timeout** field, specify the maximum time in seconds a session can remain idle before it times out and closes. The default value is 10 seconds.  
If your DNS forwarders are located at different geographical locations or if the network latency is high, you may observe session timeouts. If so, Infoblox recommends that you set the **Maximum Session Timeout** to more than 10 seconds. Increasing the session duration may impact concurrent open sessions.
5. Save the configuration.
6. As prompted, manually reboot the member to enable the DNS over HTTPS feature.



#### Note

The DNS over HTTPS feature will not take effect unless you reboot the member or the standalone system and ensure that either the DNS Cache Acceleration or Advanced DNS Protection Software service is running after the reboot.

### Configuring DNS over HTTPS in Firefox

If you are using the developer version of the Firefox browser to initiate DNS queries, you must configure additional settings in the browser to enable the DNS over HTTPS support. Complete the following steps in Firefox to enable DNS over HTTPS and upload certificates:

1. In the **Network Settings** section, click **Settings** and complete the following steps to set the Grid IP address as the custom DNS over HTTPS server:
  - a. In the *Connection Settings* dialog box select the **Enable DNS over HTTPS** checkbox.
  - b. From the **Use Provider** drop-down list, choose **Custom**.
  - c. In the **Custom** field, enter the Grid IP address in the format:  
`https://<dns-server>/dns-query`
2. Set the `network.trr.mode` preference in the configuration editor as follows:
  - a. Enter **about:config** in the Firefox address bar.
  - b. Click **Accept the Risk and Continue** to open the configuration editor.
  - c. Search for **network.trr.mode**.
  - d. Click the Edit icon and set the value to **3**.
3. If you are using a self-signed certificate, complete the following:
  - a. From the address bar, open `https://<doh_server_IP>`.
  - b. Accept the certificate.
4. If you are using a CA certificate, complete the following:
  - a. Go to **Preferences/Options** -> **Privacy and Security** -> **View Certificates** -> **Authorities** -> **Import**.
  - b. Choose the certificate.
  - c. When prompted, select the **Trust this CA to identify websites** checkbox, and restart the browser.



#### Note

For a member with the DNS Cache Acceleration service running and the DNS over HTTPS feature enabled, if you use the developer version of the Firefox browser (configured for DNS over HTTPS support) to initiate DNS queries, you must set the `network.trr.disable-ECS` preference in the configuration editor (`about:config`) to **false** for DNS data to be cached. DNS caching does not work if `network.trr.disable-ECS` is set to true.

### CLI Support for DNS over HTTPS

You can view the status of the DNS over HTTPS service, configuration, and details of active sessions using the following commands:

- `show doh-status`
- `show doh-config`

- *show doh-stats*

## Detecting and Mitigating DNS DDoS Attacks

DNS is a tempting target for attacks given that most traditional enterprise firewalls are configured to allow port 53 traffic to service DNS, which gives attackers an easy way to evade your firewall implementation. Since DNS queries are asymmetrical, they can result in a response many times larger than the query, which means that your DNS system can be used to amplify an attack.

To protect your DNS servers and DNS service performance, Infoblox provides the following features to detect and mitigate DNS DDoS attacks, such as NXDOMAIN attacks:

- Automated mitigation of phantom domain attacks (a subset of NXDOMAIN attacks), as described in [Automated Mitigation of Phantom Domain Attacks](#).
- Configurable parameters used to detect possible NXDOMAIN attacks, as described in [Detecting NXDOMAIN Attacks](#).
- Mitigation of NXDOMAIN responses by removing the LRU (Least Recently Used) items from the list of NX (non-existent) RRsets. The appliance uses the LRU list to select entries for removal from the cache when the cache utilization exceeds the allowed threshold. For more information, see [Mitigating Possible NXDOMAIN Attacks](#).
- CLI commands for configuring RRL (Response Rate Limiting) to mitigate DNS DDoS attacks by reducing the rate at which authoritative servers respond to high volumes of malicious queries, as described in [Support for RRL \(Response Rate Limiting\)](#).

Infoblox also offers the following security features to fully protect your DNS infrastructure and implementation:

- [Infoblox Advanced DNS Protection](#)
- [Infoblox DNS Firewall](#)
- [Infoblox Threat Insight](#)

## Automated Mitigation of Phantom Domain Attacks

A phantom domain attack happens when the attacker sets up "phantom" domains that do not respond to DNS queries. Under normal circumstances, the DNS recursive server contacts authoritative servers to resolve recursive queries. When phantom domain attacks happen, the recursive server continues to query non-responsive servers, which causes the recursive server to spend valuable resources waiting for responses. When resources are fully consumed, the DNS recursive server may drop legitimate queries, causing serious performance issues.

NIOS provides a few configurable parameters for mitigating phantom domain attacks in which recursive server continues to query non-responsive servers. Before you configure any of the parameters for mitigating phantom domain attacks, review the guidelines that might help you understand the relationship between these parameters. For information, see [Guidelines for Mitigating Phantom Domain Attacks](#) below.

To configure parameters for mitigating phantom domain attacks, see [Configuring Parameters for Mitigating Phantom Domain Attacks](#).

All events related to these operations are logged to the syslog. For information about the syslog and how to use it, see [Using a Syslog Server](#).

### Guidelines for Mitigating Phantom Domain Attacks

To detect phantom domain attacks, you can review your log messages. One possible indication of attack is when you receive a log message similar to the following:

```
2015-04-29T10:20:06+00:00 daemon infoblox named\[25390\]: warning no more recursive clients: quota reached
```

Consider the following guidelines to mitigate the attack:

- Increase the number of recursive clients.
- Use a combination of the following parameters to achieve optimum results:
  - **Limit recursive queries per server** and **Limit recursive queries per zone**



- **Enable holddown for non-responsive servers** and **Limit recursive queries per zone**
- When you enable any of the options, the default values are set at an optimum level for general operations. Infoblox recommends that you keep the default values when using these commands. Ensure that you understand the ramifications if you want to change the default values.

## Configuring Parameters for Mitigating Phantom Domain Attacks



### Note

Updating the parameter values for mitigating phantom domain attacks takes effect immediately through Grid replication. However, for these values to be updated in the named.conf file, you need to restart the DNS service. To restart the member service, see [Restarting Services](#).

To adjust the parameters to mitigate phantom domain attack parameters, complete the following:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox -> Edit icon. To override Grid settings, click **Override** and complete the appropriate fields.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click the **Security** tab and complete the following in the **Non-responsive servers** section:
  - **Enable holddown for non-responsive servers:** When you select this checkbox, the appliance stops sending queries to non-responsive servers for a specified time interval (**Hold down duration**) when the number of consecutive attempts to contact a non-responsive server exceeds the threshold value (**Timeouts to trigger**). No service restart is required when you change this. Clear this checkbox to disable the holddown. Note that disabling this does not clear any of the previously configured values. When you enable this feature again, the appliance preserves the previously configured values.
    - **Minimum timeout:** When the time taken for a timeout to happen exceeds this number, the timeout is counted towards the number of consecutive timeouts (**Timeouts to trigger**). You can specify a value between 0 and 5000 milliseconds. For example, if you set the minimum timeout to 1000 milliseconds, only timeouts that took longer than 1000 milliseconds are counted towards the number of consecutive timeouts. The default is 1000 milliseconds.
    - **Timeouts to trigger:** The number of consecutive timeouts before holding down a server. You can specify a value between 0 and 4294967295. For example, setting the threshold value to 5 means the appliance stops sending queries to the non-responsive server after five consecutive timeouts. The default is 5.
    - **Hold down duration:** The holddown duration for a server. You can specify a value between 1 and 86400 seconds. For example, if you set the holddown time to 60 seconds, the server stops sending queries for 60 seconds. The default is 60 seconds.  
Note in order to get enough upstream queries and for the appliance to effectively identify non-responsive servers and stop sending queries to them, do not set a high value for the **Minimum timeout** field. The higher the value you configure for this field, the longer it takes to capture a timeout and the harder it is to satisfy the total counts of consecutive timeouts (**Timeouts to trigger**). Until the total count of consecutive timeouts is exceeded, no mitigation happens against the non-responsive servers. As a result, it is less likely for the appliance to identify phantom domain attacks when you set the **Minimum timeout** field at a high value. Infoblox highly recommends that you keep the default **Minimum timeout** value to achieve optimum protection against these attacks.
  - **Limit recursive queries per server:** Select this checkbox to configure the maximum number of concurrent recursive queries that the appliance sends to a single upstream name server. Queries above the limit will be blocked and may result in a SERVFAIL response to the client. When you enable this option, the appliance dynamically adjusts the concurrent query limit for a specific server based on the ATR (Average Timeout Ratio). No service restart is required when you change this. Clear this checkbox to disable this option. Note that disabling this does not clear any of the previously configured values. When you enable this feature again, the appliance preserves the previously configured values.
    - **Maximum fetches per server:** The maximum number of concurrent recursive queries that the appliance sends to a single upstream name server before blocking additional queries to that server. You can specify a value between 0 and 4294967295. The default value is 500.



- **Quota recalculation interval:** This determines how often (in number of recursive responses) the appliance recalculates the average timeout ratio. You can specify a value between 0 and 4294967295. The default value is 200. Note that if you set this value to 0 (zero), the appliance will never recalculate the ATR. Infoblox strongly recommends that you do not set this value to 0.
- **Limit recursive queries per zone:** Select this checkbox to configure the maximum number of concurrent recursive queries the DNS server sends for a domain. If the number of recursive queries exceeds the configured value, the server blocks new queries for that domain and returns a SERVFAIL response to the client. No service restart is required when you change this. Clear this checkbox to disable the option. Note that disabling this does not clear any of the previously configured values. When you enable this feature again, the appliance preserves the previously configured values.
  - **Maximum fetches perzone:** The maximum number of concurrent recursive queries that a server sends for one of its domain. When the number of queries exceeds this number, the server blocks new queries for the domain. You can specify a value between 0 and 4294967295. The default value is 200.

## Detecting NXDOMAIN Attacks

NXDOMAIN attacks are symmetrical DDoS attacks that involve a large number of DNS clients sending queries for invalid or non-existent domains, which results in DNS recursion and NXDOMAIN responses. As a result, the DNS server spends valuable resources processing spurious requests instead of providing legitimate DNS services. When a DNS server is under NXDOMAIN attack, clients cannot get valid responses because the cache of the DNS server is flooded with NXDOMAIN results.

Infoblox provides a few options for detecting possible NXDOMAIN attacks. You can track one or all of the following to raise alerts for these attacks:

- High percentage of NXDOMAIN responses, as described in [Tracking NXDOMAIN Responses](#) below.
- Low cache hit ratio for queries, as described in [Tracking Cache Hit Ratio of Recursive Queries](#) below.
- High number of dropped UDP packets, as described in [Tracking Dropped UDP Packets](#) below.

Each of these options provides configurable parameters that determine if an alert should be raised. When an alert is triggered, the appliance sends SNMP traps about possible NXDOMAIN attacks. All triggered events are logged to the syslog. Note that you must enable notifications in order for the appliance to send SNMP traps. For more information about how to enable this, see [Setting SNMP and Email Notifications](#).



### Note

The default values for these configurable parameters are set at an optimum level for general operations. Infoblox recommends that you keep the default values when using these features. Ensure that you understand the ramifications if you want to change the default values.

## Tracking NXDOMAIN Responses

When under NXDOMAIN attack, the ratio of NXDOMAIN responses from upstream servers to all incoming recursive responses is typically high. When the ratio of NXDOMAIN responses to all incoming recursive responses exceeds the configured high water threshold, the appliance sends an alert. Note that timeouts are not counted as responses for this detection.



### Note

Changes made to the configuration for tracking NXDOMAIN responses take effect immediately on active DNS service and do not require a service restart.

To configure the parameters for tracking NXDOMAIN responses, complete the following:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox -> Edit icon.  
 To override Grid settings, click **Override** and complete the appropriate fields.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click the **Security** tab and complete the following in the **Bogus-query alerting and mitigation** section:
  - **Track the percentage of NXDOMAIN responses to recursive queries:** Select this checkbox to track the percentage of NXDOMAIN responses from up-stream servers to all incoming recursive responses. Clear this checkbox to disable the detection. Note that disabling this does not clear any of the previously configured values. When you enable this feature again, the appliance preserves the previously configured values. This is selected by default.
  - **Minimum responses per interval:** Enter the minimum number of incoming DNS responses received (within the detection interval) before the appliance starts calculating the NXDOMAIN ratio at the end of the detection interval. The appliance then compares the calculated percentage to the high water threshold. If the percentage equals or exceeds the high water threshold, the appliance sends an SNMP trap (if enabled) about possible NXDOMAIN attacks. The default is 1000. Note that the **Minimum responses per interval** is implemented to ensure that enough incoming DNS responses are received so the appliance can calculate a meaningful NXDOMAIN ratio and does not declare possible attacks from a small response sample. Therefore, when you change the default value, ensure that you use a reasonable value so the appliance calculates the NXDOMAIN ratio from a reasonable amount of responses. Raising an alert using a small response sample may not be a reliable way for detecting possible NXDOMAIN attacks.
  - **NXDOMAIN threshold:** Enter the **Low** and **High** water thresholds at which an alert is triggered. The appliance sends an alert when the percentage equals or exceeds the **High** water threshold. When the percentage in subsequent detection intervals falls below the **Low** water threshold, the appliance sends another alert to notify that the percentage of NXDOMAIN responses has gone back to an acceptable level. The defaults are 70% for **Low** and 80% for **High**.
  - **Detection interval and Responses:** These parameters work as alternatives to each other in determining when the appliance starts calculating the NXDOMAIN ratio. In the case of a very low response rate when the total responses received within the **Detection interval** never reach the **Minimum responses per interval**, the response counters continue to cumulate into subsequent detection intervals until the **Minimum responses per interval** is met. The appliance then sends an alert at the end of the detection interval. On the other hand, in the case of a very high response rate when the total number of responses received equals or exceeds the **Responses** value before the **Detection interval** is reached, the **Minimum responses per interval** does not apply and the appliance sends an alert as soon as the total responses equal or exceed the **Responses** value you define here. It does not wait till the end of the **Detection interval**. Note that the number of **Responses** you define here must be the same or greater than the **Minimum responses per interval**. The defaults are 10 seconds for the **Detection interval** and 100000 for **Responses**.
3. Save the configuration.

### Configuration Examples for Tracking NXDOMAIN Responses

The following examples demonstrate how different responses per second affect the calculation of NXDOMAIN ratio.

**Example One: Total responses per second = 250 and parameters = default values**

Detection Interval (10 seconds)											
Total Responses per Second	0	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
250/s	0	250	500	750	1000	1250	1500	1750	2000	2250	2500

In this example, the total number of responses is 250 per second and the total number of responses hits the **Minimum responses per interval** (default = 1000) at the 4th second of the detection interval. This meets the requirement for the **Minimum responses per interval** and triggers an NXDOMAIN ratio calculation at the end of the detection interval (default

= 10 seconds). If the percentage equals or exceeds the high water threshold, the appliance sends an alert and logs the event to the syslog to notify about possible NXDOMAIN flood attacks. The appliance resets the response counters for the next detection interval.

**Example Two: Total responses per second = 40 per second and parameters = default values**

1st Detection Interval (10 seconds)											
Total Responses per Second	0	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
40/s	0	40	80	120	160	200	240	280	320	360	400

2nd Detection Interval (10 seconds)											
Total Responses per Second	0	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
40/s		440	480	520	560	600	640	680	720	760	800

3rd Detection Interval (10 seconds)											
Total Responses per Second	0	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
40/s		840	880	920	960	1000	1040	1080	1120	1160	1200

In this example, the total number of responses is 40 per second. During the first detection interval of 10 seconds, the total number of responses is 400, which does not reach the **Minimum responses per interval** (default = 1000); therefore, no NXDOMAIN ratio calculation occurs and the response counters continue to accumulate into the second detection interval. At the end of the second interval, the total number of responses still does not reach the **Minimum responses per interval**; therefore, no NXDOMAIN ratio calculation occurs and the counters continue to accumulate. Finally, the total number of responses meets the requirement of the **Minimum responses per interval** when the appliance receives 1000 responses at the 5th second during the third detection interval. This triggers an NXDOMAIN ratio calculation at the end of the third detection interval, and the counters reset for the next detection interval. If the NXDOMAIN percentage equals or exceeds the high water threshold, the appliance sends an alert and logs the event to the syslog to notify about possible NXDOMAIN flood attacks.

**Example Three: Total responses per second = 50000 and parameters = default values**

Detection Interval (10 seconds)											
Total Responses per Second	0	1st	2nd	3rd	4th	5th	6th	7th	8th	9th	10th
50000	0	50000	100000								

In this example, the total number of responses per second is 50000. When the total number of responses hits 100000, which equals to the **Responses** value, the appliance starts calculating the NXDOMAIN ratio and sends an alert without waiting until the end of the detection interval. If the percentage equals or exceeds the high water threshold, the appliance

sends an alert and logs the event to the syslog to notify about possible NXDOMAIN flood attacks. The appliance resets the response counters for the next detection interval.



#### Note

The above configuration examples also apply to how the appliance tracks cache hit ratio of recursive queries, as described in Tracking Cache Hit Ratio of Recursive Queries below.

## Tracking Cache Hit Ratio of Recursive Queries

Another way to track possible DNS attacks is to monitor the cache hit ratio of recursive queries. A cache hit means the response to a query can be found in the cache of the DNS server. When the response cannot be found, it is a miss. Cache hit ratio is the percentage of cache hits to the total number of queries. The higher the ratio, the more efficiently the cache is operating. When a server is under NXDOMAIN attack, the cache hit ratio tends to drop.

To track cache hit ratio, complete the following:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox -> Edit icon. To override Grid settings, click **Override** and complete the appropriate fields.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click the **Security** tab and complete the following in the **Bogus-query alerting and mitigation** section:
  - **Track the cache hit ratio of queries:** Select this checkbox to track the cache hit ratio of recursive queries. Clear this checkbox to disable the detection. No service restart is required when you change this. Note that disabling this does not clear any of the previously configured values. When you enable this feature again, the appliance preserves the previously configured values. This is selected by default.
    - **Minimum queries per interval:** Enter the minimum number of incoming DNS queries received (within the detection interval) before the appliance starts calculating the cache hit ratio at the end of the detection interval. The appliance then compares the calculated percentage to the low water threshold. If the percentage equals or falls below the low water threshold, the appliance sends an SNMP trap (if enabled) about possible NXDOMAIN attacks. The default is 1000. Note that the **Minimum responses per interval** is implemented to ensure that enough incoming DNS queries are received so the appliance can calculate a meaningful cache hit ratio and does not declare possible attacks from a small number of incoming queries. When you change the default value, ensure that you use a number that is big enough so the appliance calculates the cache hit ratio from a reasonable amount of queries.
    - **Minimum cache utilization:** Cache hit ratio detection does not start until cache utilization hits this number. In other words, the appliance does not calculate the cache hit ratio until the cache utilization has reached or exceeded this number. The default is 75%.
    - **Hit ratio threshold:** Enter the **Low** and **High** thresholds at which an alert is triggered. The appliance sends an alert when the cache hit percentage equals or drops below the **Low** water threshold, which means the cache hit rate is low enough that the server is not operating efficiently, and there could be a high number of bogus queries that do not have matching responses in the cache. When the cache hit percentage in subsequent detection intervals reaches or exceeds the **High** water threshold, the appliance sends another alert to notify that the cache hit rate has gone back to an acceptable level. The defaults are 70% for **Low** and 80% for **High**.
    - **Hit ratio detection interval and Responses:** These parameters work as alternatives to each other in determining when the appliance starts calculating the cache hit ratio. In the case of a very low response rate when the total responses received within the **Hit ratio detection interval** never reach the **Minimum queries per interval**, the response counters continue to accumulate into subsequent detection intervals until the **Minimum queries per interval** is met. The appliance then sends an alert at the end of the detection interval. On the other hand, in the case of a very high response rate when the total number of responses received equals or exceeds the **Responses** value before the **Hit ratio detection interval** is reached, the **Minimum queries per interval** does not apply and the appliance sends an alert as soon as the total responses equal or exceed the **Responses** value you define here. It does not wait till the end of the **Hit ratio detection interval**. Note that the number of **Responses** you define here must be the same or greater than the **Minimum queries per interval**. The defaults are 10 seconds for the **Hit ratio detection interval** and 100000 for **Responses**.

3. Save the configuration.

## Tracking Dropped UDP Packets

When the DNS server starts dropping UDP packets from incoming traffic, it can be an indication of DNS attacks. Tracking dropped UDP packets can help raise awareness of possible DDoS attacks. For this feature, the appliance tracks all UDP packets, not only DNS queries. When tracking dropped UDP packets, the appliance tracks IPv4 and IPv6 packets independently.

To track dropped UDP packets, complete the following:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox -> Edit icon. To override Grid settings, click **Override** and complete the appropriate fields.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click the **Security** tab and complete the following in the **Bogus-query alerting and mitigation** section:
  - **Track the number of UDP packets dropped:** Select this checkbox to track the number of dropped UDP packets. Clear this checkbox to disable the detection. No service restart is required when you change this. Note that disabling this does not clear any of the previously configured values. When you enable this feature again, the appliance preserves the previously configured values. This is selected by default.
    - **Packet drop detection interval:** Enter the time interval of each detection period for tracking dropped UDP packets. The default is 10 seconds.
    - **Minimum packets per interval:** Enter the minimum number of UDP packets the appliance allows (within the detection interval) before it starts calculating the percentage of dropped packets to the total number of packets. The appliance compares the calculated percentage to the high water threshold. If the percentage equals or exceeds the high water threshold, the appliance sends an SNMP trap (if enabled). The default is 1000.
    - **Packet drop threshold:** Enter the **Low** and **High** water thresholds at which an alert is triggered. The appliance sends an alert when the percentage equals or exceeds the **High** water threshold. When the percentage in subsequent detection intervals equals or falls below the **Low** water threshold, the appliance sends another alert to notify that the amount of dropped UDP packets has gone back to an acceptable level. The defaults are 20% for **Low** and 30% for **High**.
3. Save the configuration.

## Mitigating Possible NXDOMAIN Attacks

To mitigate possible NXDOMAIN attacks, you can configure the appliance to split the LRU (Least Recently Used) list into two: one for NX (non-existent) RRsets and the other for all other RRsets. The LRU list is used to select entries that the appliance removes from the cache when the cache utilization exceeds the allowed threshold, which is a fraction of the configured maximum cache size.

While this option is enabled, the appliance removes the least recently used items from the LRU list for NX RRsets before removing items from the LRU list for other RRsets. Doing so helps preserve valid DNS responses in the cache while eliminating NXDOMAIN responses.

To mitigate NXDOMAIN responses, complete the following:



### Note

Changes made to the configuration for mitigating possible NXDOMAIN attacks take effect immediately on active DNS service and do not require a service restart.

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox -> Edit icon. To override Grid settings, click **Override** and complete the appropriate fields.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click the **Security** tab and complete the following in the **Bogus-query alerting and mitigation** section:

- **Deprioritize caching of NXDOMAIN responses:** Select this checkbox to split the LRU list into two: one for NX RRsets and the other for all other RRsets, and to always remove the least recently used items from the list of NX RRsets first. This is selected by default.
3. Save the configuration.

## Support for RRL (Response Rate Limiting)

RRL provides the ability to control excessive UDP responses that are identical or almost identical. For more information about how RRL works, refer to the *BIND Administrators Reference Manual* at <http://www.isc.org/downloads/bind/doc/>.

You can configure parameters for this feature through the following CLI commands: `set dns_rrl` and `show dns_rrl`. For more information about these commands, refer to the Infoblox CLI Guide.

You can also log RRL events to the syslog. To enable RRL logging, select **rate-limit** in the **Logging Category** when you configure logging for the Grid or member. For more information about how to select logging categories, see [Setting DNS Logging Categories](#).

## Blacklists

Your organization can prevent customers or employees from accessing certain Internet resources, particularly web sites, by prohibiting a recursive DNS member from resolving queries for domain names that you specify.

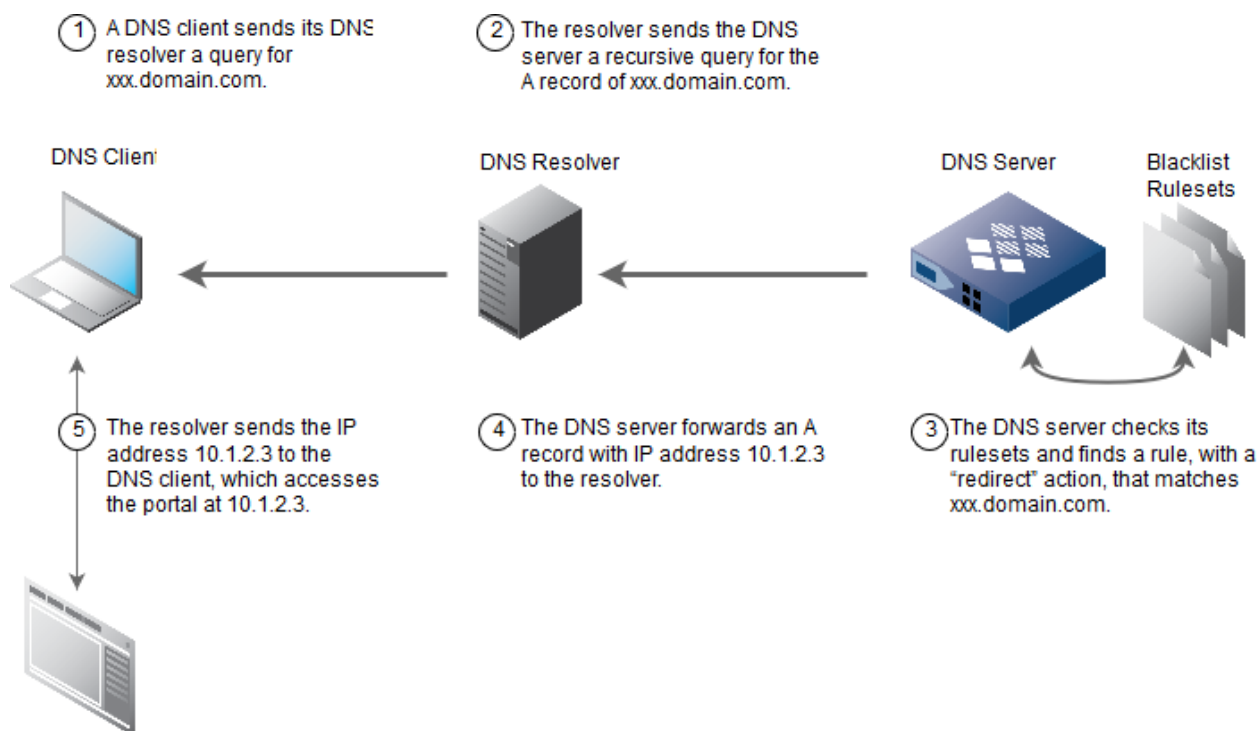
You can create blacklist rules that specify how a DNS member responds to recursive queries for data for which it is not authoritative. Each rule specifies a domain name and the action of the DNS member when the domain name in the query matches that in the rule. Instead of resolving the query, the DNS member can redirect the DNS client to predefined IP addresses or return a REFUSED response code indicating that resolution is not performed because of local policy.

When the DNS member receives a query for data for which it is not authoritative, it first tries to match the domain name in the query with a domain name in any of its rules. If it finds a match, it responds according to the action specified in the rule. If it does not find a match and the NXDOMAIN feature is enabled, the DNS member checks the NXDOMAIN rulesets for a match and responds accordingly. If the NXDOMAIN feature is not enabled, the DNS member resolves the query. (For information about the NXDOMAIN feature, see [About NXDOMAIN Redirection](#).)

Infoblox DNS members can modify their responses to queries for A records only. Therefore, if the matched query is for a record other than an A record, including a query with a type of "ANY", the DNS member sends a REFUSED response if the matched rule has an action of "Redirect".

In the figure Blacklist below, a DNS client opens a web browser and tries to access xxx.domain.com. When the DNS member receives the query for xxx.domain.com, it checks its blacklist rulesets and finds xxx.domain.com in a rule with an action of "Redirect". The DNS client is redirected to the configured redirection destination IP address 10.1.2.3.

*Blacklist*



This feature supports queries for data in IPv4 and IPv6 reverse-mapping zones, as well as forward-mapping zones. Note that when a user with a Windows DNS client with IPv6 installed tries to access a domain name, the Windows client sends queries for AAAA records before queries for A records. After the DNS member sends a Refused response to the query for the AAAA record, the DNS client then sends a query for the A record. The DNS member then responds according to the blacklist rules.

When DNSSEC is enabled on the Infoblox DNS server, it does not redirect DNS clients that request DNSSEC data. (For information about DNSSEC, see [Configuring DNSSEC](#).) If DNSSEC is not enabled and the query includes a request for DNS data, the appliance ignores the request for DNSSEC data and redirects the clients.

To apply the configured DNS blacklist rules regardless of whether a DNS query requests DNSSEC data, configure the appliance accordingly. For more information about how to configure this, see [Applying Policies and Rules to DNS Queries that Request DNSSEC Data](#).

You can enable the blacklist feature at the Grid, member, and DNS view levels. Note that only recursive DNS servers can support this feature. For information on enabling recursion on a DNS member, see [Enabling Recursive Queries](#).

## About Blacklist Rulesets

A blacklist ruleset is a list of rules that a DNS member uses to determine its response to recursive queries for certain domain names. When you enable the blacklist feature, you must define at least one rule in a ruleset. Each rule consists of a domain name and an associated action. The DNS member matches the domain names in the rules with the entire domain name in the query, including its suffix. The domain name in the rule can contain any printable character. Domain name matching is case-insensitive. Unlike the NXDOMAIN rules, blacklist rules do not support metacharacters in domain names.

The action in a rule is either "Pass" or "Redirect".

- Pass: The DNS member resolves the query and forwards the response to the DNS client.
- Redirect: The DNS member does not resolve the query. The DNS member redirects the client to the predefined IP addresses or sends a REFUSED response, depending on your configuration. Note that the DNS member can redirect the client only if the query is for an A record. If the query is for another resource record, the DNS member sends a REFUSED response.

You can use the Blacklist wizard, described in [Adding a Blacklist Ruleset](#), to add blacklist rulesets, but not rules. You can only add rules by importing them in a CSV file, as described in [Importing and Exporting Data using CSV Import](#). Note that if a blacklist ruleset contains duplicate domain names, the DNS member loads the first rule in the ruleset and discards the other rules.

The following example illustrates how the DNS member applies blacklist rules. Ruleset 1:

Pattern	Action
a1.foo.com	PASS
foo.com	REDIRECT/BLOCK

- If the DNS member receives a recursive query for a1.foo.com, it resolves the query and forwards the response to the client.
- If the DNS member receives a recursive query for the A record of b1.foo.com, it redirects the DNS client to the specified IP address. If the query is for another record type, such as an MX record, the member sends a REFUSED response to the client.

## Blacklist Guidelines

The following summarizes how a DNS member responds to a DNS client when the blacklist feature is enabled:

- If the domain name in the query matches a domain name in a rule, the member does the following:
  - If the query is for an A record, the member performs the action specified in the rule.
    - If the action is "Redirect", the member performs the action specified in the Blacklist wizard.
      - If the action in the wizard is to redirect, the DNS member redirects the client to the listed IP addresses.
      - If the action in the wizard is to return a REFUSED response, the DNS member sends a REFUSED response to the DNS client.
    - If the action in the rule is "Pass", the DNS member resolves the query and forwards the response to the DNS client.
  - If the query is for a non-A record, the member performs the action in the rule as follows:
    - If the action is "Redirect", the DNS member returns a REFUSED response to the DNS client.
    - If the action is "Pass", the DNS member resolves the query and forwards the response to the DNS client.
- If the domain name in the query does not match a domain name in a rule:
  - If the NXDOMAIN feature is enabled, the DNS member tries to find a match with the NXDOMAIN rules and responds accordingly.
  - If the NXDOMAIN feature is disabled, the DNS member resolves the query and forwards the response to the DNS client. Note that if an A record with a dotted hostname is added to an authoritative zone through a dynamic DNS update, and that A record should actually belong in an existing delegation, the appliance may not redirect a query for that A record according to the Blacklist and NXDOMAIN guidelines.

Related topic

[Configuring Blacklists](#)

## Configuring Blacklists

You can perform the following operations on the blacklist feature:

1. Add blacklist rulesets, as described in [Adding a Blacklist Ruleset](#) below.
2. Create one or more CSV files that contain the rules for each ruleset and import the files to the Grid. For information about importing CSV files, see [Importing and Exporting Data using CSV Import](#).
3. Enable blacklisting, as described in [Enabling Blacklisting](#) below.



## Adding a Blacklist Ruleset

To add the name of a blacklist ruleset:

1. From the **Data Management** tab -> **DNS** tab -> **Blacklist Rulesets** tab, click the Add icon.
2. In the *Blacklist* wizard, complete the following:
  - **Name:** Enter a name for the ruleset.
  - **Comment:** You can enter additional information.
  - **Disable:** You can disable this ruleset for use later on. The appliance ignores disabled rulesets.
3. Save the configuration and click **Restart** if it appears at the top of the screen. You can then use the CSV Import feature to import the rules for each ruleset.

## Managing Blacklist Rulesets

To view rulesets, navigate to the **Data Management** tab -> **DNS** tab -> **Blacklist Rulesets** tab. The panel lists the configured rulesets and their associated comments. You can also display the Disabled column which indicates which rulesets are disabled. From this panel, you can do the following:

- Add more rulesets, as described in the preceding section, Adding a Blacklist Ruleset.
- Edit a ruleset, by clicking its checkbox and clicking the Edit icon. You can set the following in the Blacklist Ruleset editor:
  - In the **General Basic** tab, you can change entries in any of the fields.
  - In the **Permissions** tab, you can set admin permissions for the ruleset.
- Delete a ruleset, by clicking its checkbox and clicking the Delete icon.
- View the rules that were imported in each ruleset by selecting it. For each rule, the panel displays the following:
  - Domain name
  - The action of the recursive DNS member when the domain name in a query matches the domain name in the rule.

To delete or edit rules in a ruleset, you must delete the ruleset from this panel, edit the CSV file and re-import it.

## Enabling Blacklisting

Only DNS members with recursion enabled can support this feature. You can enable this feature at the Grid level and override it for a member or DNS view with recursion enabled.

You can also enable the DNS member to log queries that matched blacklist rules. The logs include the queried domain name, source IP address, the pattern of the matched rule, and the name of the corresponding ruleset.

To enable blacklisting:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
**DNS View:** From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns\_view* checkbox -> Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. If the *Grid DNS Properties* or *Member DNS Properties* editor is in basic mode, click **Toggle Advanced Mode**.
3. Click **Blacklist** and complete the following:
  - Enable Domain Name Blacklist:** Select this checkbox.
  - Blacklist Rulesets:** To add a ruleset, click the Add icon. If there are multiple rulesets, select one from the *Select Ruleset* dialog box. Use the up and down arrows to move rulesets up and down in the list. The appliance applies rulesets in the order they are listed.
  - For blacklisted domain names, return:** Select the action of the appliance when it receives a query for a record that matches a rule with an action of Redirect/Block.  
If you selected **This list of IP addresses**, add an IP address to the **Redirect to** table by clicking the Add icon and entering the address. The addresses are listed in round robin fashion in the synthesized response of the DNS member. You can enter up to 12 IP addresses.
  - Blacklist TTL:** Specify how long the DNS client caches the A record with the redirected IP address.
  - Log queries for blacklisted domain names:** Select this option to enable the appliance to log queries for blacklisted domain names, including the source IP address of the query.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

Related topic

[Blacklists](#)

## About Root Name Servers

Root name servers contain the root zone file which lists the names and IP addresses of the authoritative name servers for each top-level zone. When a root name server receives a query for a domain name, it provides at least the names and addresses of the name servers that are authoritative for the top-level zone of the domain name.

You can configure the NIOS appliance to use Internet root name servers or custom root name servers. If you enable recursive queries and the appliance receives a recursive query that it cannot resolve locally, then it queries specified forwarders (if any) and then queries any root name servers that you configure. If you do not specify internal root name servers and the appliance can access the Internet, it queries the Internet root name servers.

You can specify root name servers for the Grid, individual members, and user-defined DNS views. You can specify root name servers for all DNS views except the default view. The default view uses either the member level root name servers (if specified) or the Grid level root name servers.

Every Grid member has a default view. If you want to specify root name servers for a default view, override the Grid root name server setting at the member level and the default view can use the member-level setting.

## Specifying Root Name Servers

To specify root name servers for a Grid, member, or DNS view:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab, click the **Members** tab -> *member* checkbox -> Edit icon.  
**DNS View:** From the **Data Management** tab, select the **DNS** tab, click the **Zones** tab-> *dns\_view* checkbox -> Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the *Grid DNS Properties* and *Member DNS Properties* editors, you must click **Toggle Advanced Mode**.
3. When the additional tabs appear, click **Root Name Servers**.
4. Select one of the following options:
  - Use Internet root name servers: This option is selected by default.
  - **Use custom root name servers:**
    - Click the Add icon and enter the following information when a new row appears:
      - **Name:** Enter a name for the root name server.
      - **Address:** Enter the IP address of the root name server. The feature supports IPv4 or IPv6 values.
    - Optionally, you can perform the following:
      - Select a server from the custom root name servers list, and then click the Edit icon to modify its information.
      - Select a server from the custom root name servers list, and then click the Delete icon.
    - *Member DNS Properties* editor only: You can choose whether the DNS views should inherit the custom root name server from the member. Select one of the following:
      - **Applies to default DNS view only:** Select this option to apply the root name server only to the default DNS view on the selected member. This option is selected by default.
      - **Applies to all DNS Views on this member:** Select this option to apply the root name server for all DNS views on the selected member.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

## About Sort Lists

A sort list prioritizes A and AAAA records on certain networks when those records are included in responses, sorting them to the beginning of the list in the response. For example, you can define a sort list when a server has two interfaces and

you want the DNS clients to prefer one interface because it has a faster link. When you define a sort list on the NIOS appliance, you specify the following:

- The IP address or network of the source of the query
- The IP addresses or networks that the appliance lists first in its response when it receives a query from the corresponding source address

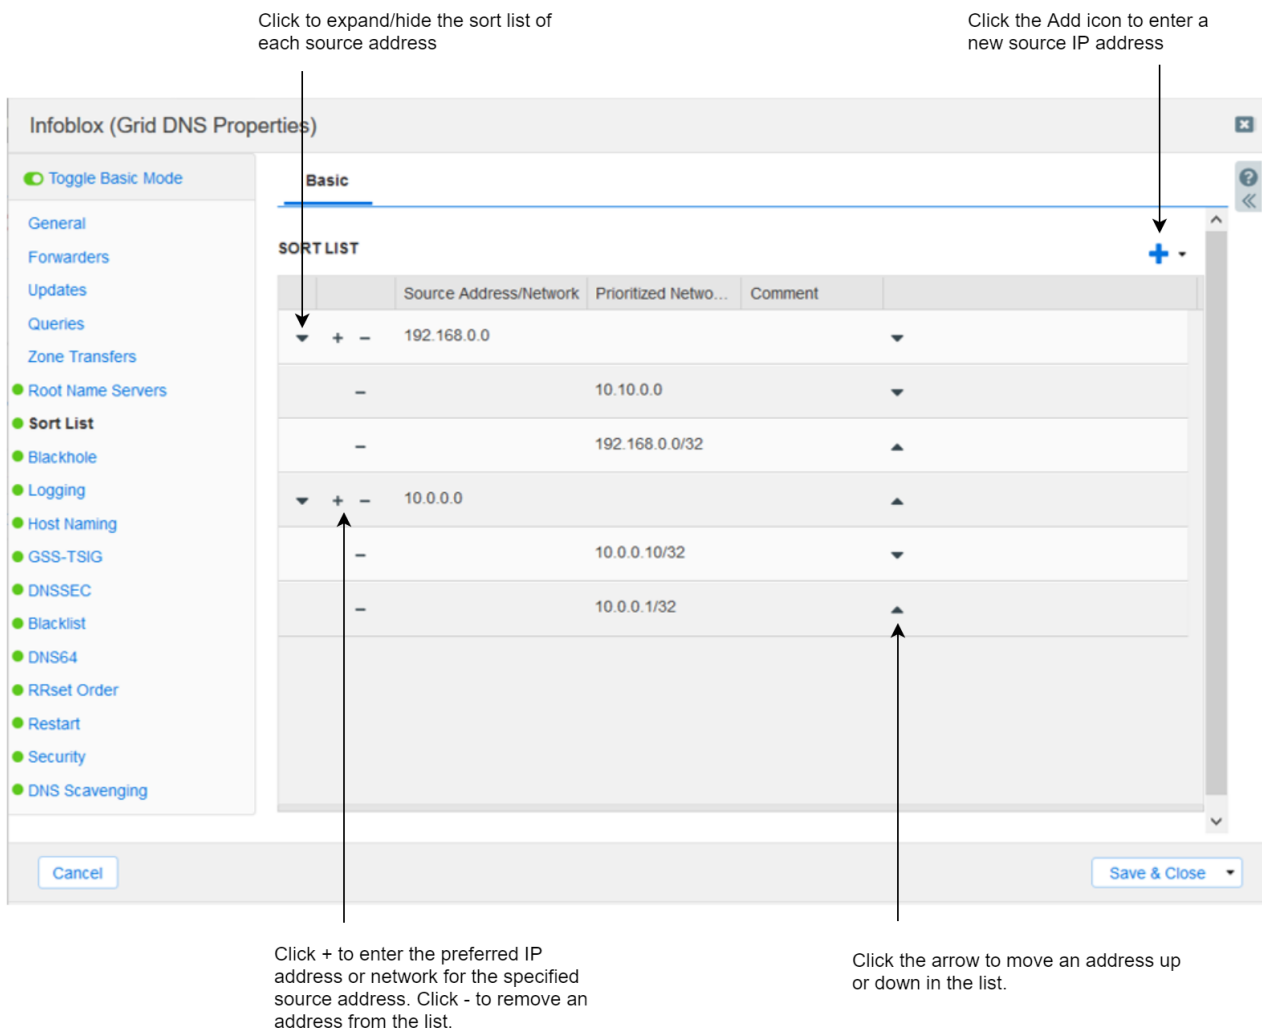
When the NIOS appliance receives a query from the specified IP address or network and the DNS lookup produces a response with multiple addresses, the NIOS appliance sorts the addresses so that those in the sort list are at the beginning of its response. For more information, refer to the *Infoblox DNS Cache Acceleration Application Guide*.

## Defining a Sort List

To define a sort list for a Grid, member, or DNS view:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab > **Members** tab -> *member* checkbox -> Edit icon.  
**DNS View:** From the **Data Management** tab, click the **DNS** tab -> **Zones** tab > *dns\_view* checkbox -> Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click **Sort List**.
4. Click the Add icon and select either Any to define a sort list for any address and network, or Address/Network to define a sort list for a particular source IP address or network.
5. Do the following in the new row:
  - If you selected Address/Network, enter the IP address or network of the source of the query. The feature supports IPv4 or IPv6 values.
  - Click the Add icon beside the source IP address to add the preferred IP addresses or networks for the source. You can add as many IP addresses as necessary. When you add multiple IP addresses, you can change the order of the IP addresses. Select an IP address and drag it to its new position, or click the up or down arrow, as shown in the figure Sort List below.

Sort List



## Configuring a DNS Blackhole List

The DNS blackhole feature provides the ability to specify IP and network addresses of network devices that you do not want to use in the DNS resolution process. The DNS blackhole feature is disabled by default. When enabled, the NIOS appliance does not accept queries from IP addresses in the blackhole list and does not use them to resolve queries. For example, you can add the IP addresses of name servers that are using DNS incorrectly to prevent the NIOS appliance from accepting their queries and from using them as resolvers. You can also use this feature to fix temporary network issues. For example, you can add the IP addresses of delegated servers, configured forwarders, and DHCP servers that have temporary DNS-related issues.

You can create a DNS blackhole list for the entire Grid or create a separate list for each Grid member. For example, if one of your Grid members is behind a firewall, you might need to configure a different DNS blackhole list for this member because the clients that can access it might be mapped differently.

The appliance accepts queries from addresses and networks that are excluded from the blackhole list and uses these addresses and networks as resolvers. To add an IP address to the blackhole list, enter it and set its permission to **Include**. You can also add an IP address to the blackhole list and set its permission to **Exclude** so its not in the blackhole list, effectively allowing the NIOS appliance to respond to queries from that address and to use it as a resolver.

When you add a network to a DNS blackhole list, all the IP addresses in the network are not used in the DNS resolution process. If you want to allow some IP addresses within the network, add these addresses to the list and set their permission to "Exclude." Ensure that you list these IP addresses before the network address because the appliance applies permissions to the addresses in the order they are listed. For example, when you add the network 10.10.0.0/24 to a DNS blackhole list, all 256 IP addresses in the network are put on the blackhole list. To allow DNS traffic to the IP addresses 10.10.0.55 and 10.10.0.88, add these two addresses before the network address in the DNS blackhole list,

and then set their permissions to **Exclude**.

You can define ACEs or a named ACL to determine the IPv4 and IPv6 addresses and networks that you want to include in or exclude from a blackhole list.

## Defining a DNS Blackhole List

To enable the DNS blackhole feature and configure a DNS blackhole list for a Grid or member:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab, click the **Members** tab -> *member* checkbox -> Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the editor, click the **Blackhole** tab and complete the following:  
**Enable Blackhole:** Select this checkbox to enable the DNS blackhole feature. This is disabled by default.
3. Select one of the following:
  - **Any:** Select this if you want to configure a blackhole list. The appliance will allow all clients to resolve DNS queries. This is selected by default.
  - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this, the appliance uses clients that have the **Exclude** permission in the DNS resolution process. You can click **Clear** to remove the selected named ACL.
  - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
    - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address of the client. The **Permission** column displays **Include** by default. You can change it to **Exclude** by clicking the field and selecting **Exclude** from the drop-down list. When you select **Include**, the appliance adds the IP address to the blackhole list and does not allow DNS queries and DNS resolution for this address. When you select **Exclude**, the appliance excludes the address from the blackhole list and allows DNS queries and resolution for the address.
    - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
      - **Permission:** Select **Allow** or **Deny** from the drop-down list.
    - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
      - **Permission:** Select **Allow** or **Deny** from the drop-down list.
    - **Any Address/Network:** Select to include or exclude any IP addresses and networks for the DNS resolution process.
  - After you have added access control entries, you can do the following:
    - Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
    - Reorder the list of ACEs using the up and down arrows next to the table.
    - Select an ACE and click the Edit icon to modify the entry.
    - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
4. Save the configuration.

## Specifying Hostname Policies

You can enforce a naming policy for the hostnames of A, AAAA, Host, MX, NS, and bulk host records based on user-defined or default patterns. For MX and NS records, the hostname restrictions apply to the text in the RDATA field (right-hand side) of the resource record name.

Records that you created before you enabled the hostname checking policy need not comply with the hostname restriction that you specify.

You can select one of three preconfigured policies or define your own host naming policy with a POSIX regular expression. The policies Infoblox provides implement standard host naming restrictions according to *RFC 952, DOD Internet Host Table Specification*, and *RFC 1123, Requirements for Internet Hosts -- Application and Support*.



### Note

The hostname restriction limits the hostname of A, AAAA, Host, MX, NS, and bulk host records only.

You can define your own hostname restriction policy at the Grid level only. At the member and zone levels, you can select a predefined policy or a policy that was defined at the Grid level. The appliance supports IDNs for DNS zones and resource records. For more information about IDNs, see [Managing Internationalized Domain Names](#). You can use UTF-8 characters when you configure your own hostname checking policy.

## Defining Grid Hostname Policies

You can define new hostname policies and set the hostname policy for all zones in the Grid as follows:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the *Grid DNS Properties* editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click **Host Naming**.

The Host Name Policies section lists the following preconfigured record policies:

- **Strict Hostname Checking:** You can only use hostnames that contain alphanumeric characters, dashes (-), and asterisks (\*). You cannot use other special characters, such as underscore (\_). Note that when you select this policy, the appliance automatically applies the policy to dynamic DNS updates and zone transfers it receives. When you select this, you can enter host names through Grid Manager using punycode, but not IDNs. The appliance stores IDNs that are created through DDNS updates and DNS transfers in punycode. You can monitor non-compliant host names using the Hostname Compliance report. For information, see [Obtaining a List of Invalid Record Names](#) below.
- **Allow Underscore:** You can only use hostnames with alphanumeric characters, dashes, and underscores ("-\_" and "\_"). This is the default.
- **Allow Any:** You can use any hostname.

Select **Default** from the drop-down list in the **Default** column to change the Grid default hostname policy.

From NIOS 8.6.3 onwards, you have to manually clear the existing host name policy in the table and enter the new one.

4. Click **Add** to define your own hostname checking policy.  
Enter a record policy name and a regular expression string, and click **OK**. See [Supported Expressions for Search Parameters](#) for definitions of regular expressions.  
Note that Grid Manager does not validate the regular expressions that you enter. Therefore, you can inadvertently specify an invalid regular expression that might cause noncompliance errors when you create records.
5. If you select the Strict Hostname Checking policy, the **Apply policy to dynamic updates and inbound zone transfers (requires Strict Hostname Checking setting)** option is enabled by default. It enables the appliance to apply the policy to dynamic DNS updates and zone transfers that it receives. You can then select which action the appliance takes when it encounters names that do not conform to the policy. Select either **Fail** or **Warn**. If you select **Warn**, the appliance allows the dynamic DNS update or zone transfer, but logs a syslog message.  
Note that the Strict Hostname Checking policy only allows alphanumeric characters and dashes ("-"). In addition, this policy allows IDNs that are written in punycode. You cannot use other special characters, such as underscore ("\_"). Therefore, DDNS updates from Microsoft Active Directory controllers may not be accepted.
6. Save the configuration and click **Restart** if it appears at the top of the screen.

After you specify a hostname restriction policy, if you create a record name that does not comply with this policy and try to save it, an error message appears.

## Defining Hostname Restrictions

You can select a hostname restriction policy for an individual Grid member or zone. You can specify hostname restrictions for authoritative forward-mapping zones only. You cannot specify hostname restrictions for forward zones, stub zones, IPv4 reverse-mapping zones, and IPv6 reverse mapping zones.

To select a hostname restriction policy for a Grid member or zone:

1. **Member:** From the **Data Management** tab, select the **DNS** tab, click the **Members** tab -> *member* checkbox -> Edit icon.  
**Zone:** From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab-> *dns\_view* -> *zone* checkbox -> Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the *Member DNS Properties* editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click **Host Naming**.
4. Click **Override**.
5. From the **Host Name Policy** drop-down list, select a predefined policy or a policy that was defined at the Grid level.
6. If you select the Strict Hostname Checking policy, the Apply policy to dynamic updates and inbound zone transfers (requires Strict Hostname Checking setting) is enabled by default. It enables the appliance to apply the policy to dynamic DNS updates and zone transfers that it receives. You can then select which action the appliance takes when it encounters names that do not conform to the policy. Select either Fail or Warn. If you select Warn, the appliance allows the dynamic DNS update or zone transfer, but logs a syslog message. Note that the strict hostname checking policy only allows alphanumeric characters and dashes. It does not allow for the use of other special characters, such as underscore ("\_"). Therefore, DDNS updates from Microsoft Active Directory controllers might not be accepted.
7. Save the configuration and click **Restart** if it appears at the top of the screen.

## Obtaining a List of Invalid Record Names

You can retrieve a list of all record names that do not comply with the current hostname checking policy of a zone. These could be records that were created before the current host naming policy was set. In addition, if you selected the Strict Hostname Checking policy and allowed illegal hostnames in DDNS updates and inbound zone transfers with a warning, those records are listed in this report as well.

To display the Hostname Compliance report:

1. From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab-> *dns\_view* -> *zone* checkbox.
2. Click **Hostname Compliance**.  
The Hostname Compliance Report for the zone displays. It lists the record name, type, value, and comment for all records that do not comply with the hostname restriction policy of the zone.

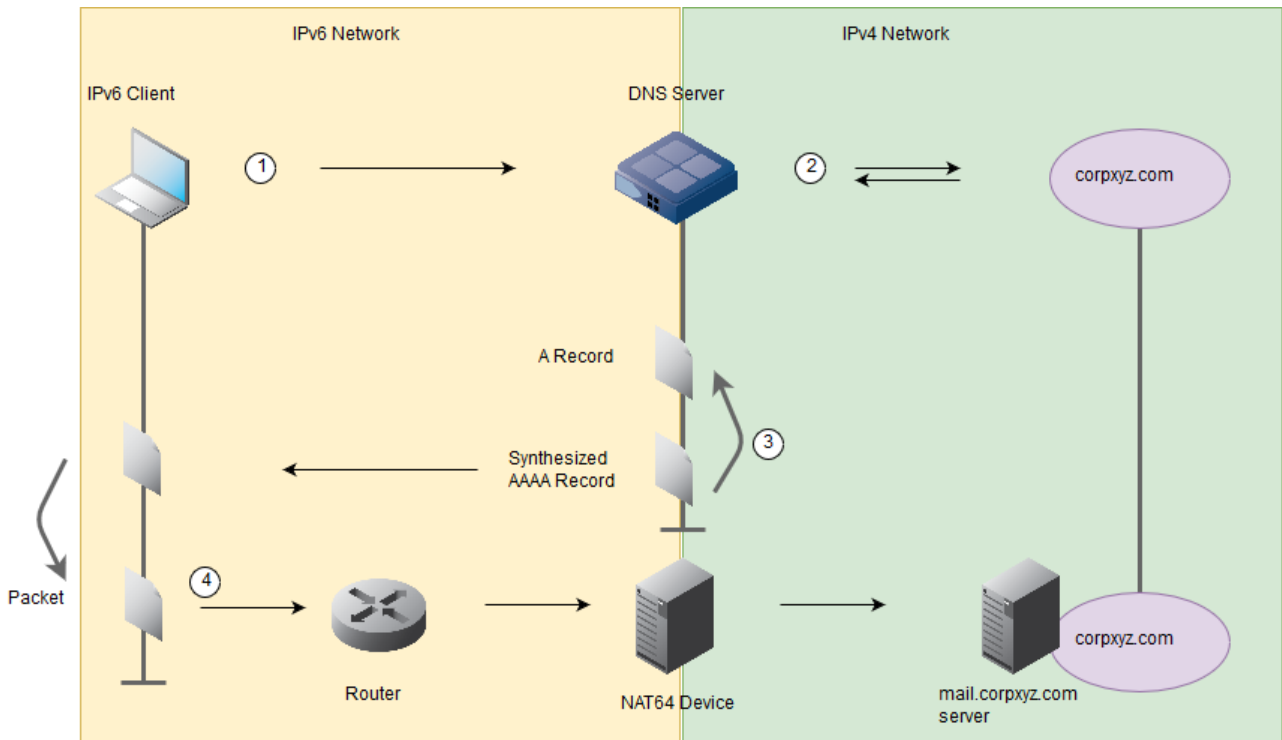
From the report, you can select a record and do the following:

- Click the Edit icon to open the record editor.
- Click the Delete icon to move it to the Recycle Bin.

## About DNS64

To support the increasing number of IPv6 and dual-stack networks, Infoblox DNS servers now support DNS64, a mechanism that synthesizes AAAA records from A records when no AAAA records exist. When you enable DNS64 on an Infoblox DNS server, it can operate with a third-party NAT64 device so IPv6-only nodes can communicate with IPv4-only nodes without any changes to either of the devices.

As illustrated in the following figure, when an IPv6-only host requests the AAAA record of an IPv4-only server and none exists, a DNS64-enabled server can retrieve the A record of the IPv4 server and synthesize an AAAA record. The IPv6-only host can then use the synthesized AAAA record, which contains the IPv6 proxy address for the IPv4 address in the original A record, to initiate communication with the IPv4 host.



Following are the steps illustrated in the above figure:

1. An IPv6-only host sends a recursive query for the AAAA record of the IPv4 server mail1.corpxyz.com.
2. The Infoblox DNS server attempts to resolve the request for the AAAA record, and determines that an AAAA record for mail1.corpxyz.com does not exist. The DNS server then performs a query for the A record of mail1.corpxyz.com.
3. The DNS server creates a synthetic AAAA resource record from the information in the A record, and returns the synthesized AAAA record to the requesting IPv6 host.
4. The host receives the synthetic AAAA record and sends a packet to the destination address specified in the synthetic AAAA record. The packet is routed to the IPv6 interface of the NAT64 device, which translates the packet from IPv6 to IPv4 and forwards it to the server, mail1.corpxyz.com.

Infoblox DNS servers can return synthesized AAAA records to both IPv4 and IPv6 clients when the client explicitly requests an AAAA record and none exists for the requested host. If a host has multiple A records, the DNS server synthesizes an AAAA record for each A record.

Infoblox DNS servers can also synthesize records for reverse-mapping zones. When a DNS server receives a query for a PTR record in the IP6.ARPA domain whose address matches a configured DNS64 prefix, the server synthesizes a CNAME record that contains an IPv4 address derived from the IPv6 address in the query. The server then sends a query for the PTR record so it can resolve the IPv4 address to the hostname.

For example, if a DNS server that is configured to synthesize records for the prefix 2001:db8::/96 receives a query for the PTR record of 2001:db8::0102:0304, it synthesizes a CNAME record that contains the IPv4 address 4.3.2.1.in-addr.arpa. The server then resolves the PTR record of the IPv4 address 4.3.2.1.in-addr.arpa.

If the server obtains the PTR record, then it sends the synthesized CNAME record and the PTR record to the client. If the zone exists, but there is no PTR record, then the server sends the synthesized CNAME record only. If the zone does not exist, then the server responds with a SERVFAIL with no answers.

Additionally, Infoblox DNS servers can generate synthesized records for DNSSEC secure zones, but only for non-DNSSEC clients. A DNS client or resolver includes the EDNS OPT pseudo-RR with the DO (DNSSEC OK) bit set to indicate that they are requesting DNSSEC data. DNS servers can generate synthesized AAAA records only when the request does not have the DO bit set. This ensures that DNSSEC clients receive only valid responses.

For additional information about DNS64, refer to the following Internet drafts:

- <http://tools.ietf.org/html/draft-ietf-behave-dns64-11>
- <http://tools.ietf.org/html/draft-ietf-behave-address-format-10>



## Configuring DNS64

You can enable DNS64 on both authoritative and recursive DNS servers. You can configure DNS64 at the Grid, member or DNS view level.

To configure DNS64 on Infoblox DNS servers:

1. Create at least one DNS64 synthesis group. A synthesis group specifies the IPv6 prefix of the synthesized AAAA records. For more information, see [Adding a DNS64 Synthesis Group](#) below.
2. Optionally, specify additional parameters for the synthesis group. For more information, see [Setting DNS64 Group Properties](#) below.
3. Enable the DNS64 service and assign a synthesis group to the Grid, a member or a DNS view. For more information, see [Enabling DNS64 Service](#) below.

On the NAT64 device, you must specify the IPv6 prefixes that are configured on the DNS server.

## About Synthesis Groups

A synthesis group specifies, among other things, the IPv6 prefix for the synthesized AAAA records. Infoblox DNS servers provide a default DNS64 synthesis group with the well-known prefix 64:ff9b::/96, which is reserved for representing IPv4 addresses in the IPv6 address space. You can keep the default group, change the prefix or delete the group. You can also add a synthesis group for a Network-Specific Prefix (NSP), which is an IPv6 prefix assigned to an organization to create IPv6 representations of IPv4 addresses.

After you create a synthesis group, you can define rules to restrict the synthesis of AAAA records to certain IPv4 addresses and networks, and specify the DNS clients and networks to which the server can send synthesized AAAA records. For more information, see [Setting DNS64 Group Properties](#) below.

Note that though you can control the synthesis of AAAA records, the DNS server always synthesizes CNAME records when it receives a query for an IPv6 PTR record whose address matches a prefix in a DNS64 synthesis group. You can also configure the DNS server to generate synthesized AAAA records for DNS queries that have the DO bit set.

## Adding a DNS64 Synthesis Group

To add a synthesis group:

1. From the **Data Management** tab, select the **DNS** tab -> **DNS64 Groups** tab, and then click the Add icon.
2. In the *DNS64 Synthesis Group* wizard, complete the following:
  - **Name:** Enter a name for the group.
  - **Prefix:** The IPv6 prefix used for the synthesized AAAA records. The default is the well-known prefix 64:FF9B::/96. The prefix length must be /32, /40, /48, /56, /64, and /96, and all bits beyond the specified length must be zero.
  - **Comment:** Optionally, enter additional information about the group.
  - **Disabled:** Select this checkbox if you would like to disable the group at this time. Note that you cannot disable the group if it is the only group that is used by a Grid, member or DNS view that has DNS64 enabled.
  - **Apply to queries requesting DNSSEC records:** Select this to generate synthesized AAAA records for DNS64 synthesis groups that request DNSSEC data.
3. Click **Next** to define extensible attributes for the synthesis group. For information, see [Using Extensible Attributes](#).
4. Save the configuration

## Viewing DNS64 Synthesis Groups

To view synthesis groups, from the **Data Management** tab, select the **DNS** tab -> **DNS64 Groups** tab. This tab displays the following information about each group:

- **Name:** The group name.
- **Prefix:** The IPv6 prefix that is assigned to the group.
- **Comment:** The comment that was entered for the group.
- **Site:** The value of this attribute, if specified.

You can display the following additional column:

- **Disabled:** Indicates whether the group is disabled.

You can do the following:

- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).
- Edit the properties of a synthesis group.
  - Select the synthesis group, and then click the Edit icon.
- Move a synthesis group to the Recycle Bin.
  - Select the synthesis group, and then click the Delete icon. Note that you cannot delete a synthesis group that is assigned to a Grid, member or DNS view.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filter](#).
- Export the synthesis groups to a .csv file.
  - Click the Export icon.
- Print the list of synthesis groups.
  - Click the Print icon.

## Setting DNS64 Group Properties

After you create a DNS64 synthesis group, you can specify the following:

- The IPv4 and IPv6 DNS clients and networks to which the DNS server is allowed to send synthesized AAAA records with the specified IPv6 prefix.
- The IPv4 addresses and networks for which the DNS server can synthesize AAAA records with the specified prefix.
- IPv6 addresses or prefix ranges that cannot be used by IPv6 only hosts, such as IP addresses in the ::ffff:0:0/96 network. When the DNS server retrieves an AAAA record that contains an IPv6 address that matches an excluded address, it does not return the AAAA record. Instead, it synthesizes an AAAA record from the A record. Note that a DNS server synthesizes the AAAA record of a host that has both A and AAAA records when all the IPv6 addresses in the AAAA records match the excluded addresses. If the host has multiple AAAA records and some of them contain excluded IPv6 addresses, then the server returns the remaining AAAA records.

You can add individual access control entries (ACEs) or use a named access control list (ACL) to define these clients.

For information about how to define named ACLs, see [Defining Named ACLs](#).

To configure DNS64 group properties:

1. From the **Data Management** tab, select the **DNS** tab -> **DNS64 Groups** tab -> *group* checkbox -> Edit icon.
2. In the **General** tab of the *DNS64 Synthesis Groups* editor, you can do the following:

- Modify the name, prefix or comment.
- Select the **Disabled** checkbox, if you want to disable the group at this time.
- Select the **Apply to queries requesting DNSSEC records** checkbox to have the DNS server generate synthesized AAAA records for DNS64 synthesis groups that request DNSSEC data.

**Perform DNS64 synthesis for these clients:** Specify IPv4 and IPv6 hosts and networks to which Infoblox DNS servers can send synthesized AAAA records. The default is to allow any IPv4 and IPv6 address and network. Select one of the following:

- **None:** Select this if you do not want to define specific addresses or networks to which the appliance sends synthesized AAAA records. When you select this, the appliance sends synthesized AAAA records to all clients. This is selected by default.
- **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance sends synthesized AAAA records to the clients that have the **Allow** permission in the list. You can click **Clear** to remove the selected named ACL.
- **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the

selected item or expands the panel so you can specify additional information about the item you are adding, as follows.

- **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
- **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
  - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
  - **Permission:** Select **Allow** or **Deny** from the drop-down list.
- **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
  - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
  - **Permission:** Select **Allow** or **Deny** from the drop-down list.
- **Any Address/Network:** Select this to allow or deny any IP addresses to which the appliance sends synthesized AAAA records.

**Mapped IPv4 Addresses:** Specify IPv4 addresses and networks for which the DNS server synthesizes AAAA records. The default is to allow the DNS server to synthesize AAAA records for any IPv4 address in any network. Select one of the following1:

- **None:** Select this if you do not want to define specific IPv4 addresses or networks for which the DNS server synthesizes AAAA records. The appliance synthesizes AAAA records for all IPv4 clients. This is selected by default.
- **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance synthesizes AAAA records for the clients that have the **Allow** permission in the list. You can click **Clear** to remove the selected named ACL.
- **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
- **IPv4 Address:** Select this to add an IPv4 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
- **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
  - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
  - **Permission:** Select **Allow** or **Deny** from the drop-down list.
- **Any Address/Network:** Select this to allow or deny any IPv4 addresses for which the appliance synthesizes AAAA records.

After you have added access control entries, you can do the following:

- Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
- Reorder the list of ACEs using the up and down arrows next to the table.
- Select an ACE and click the Edit icon to modify the entry.
- Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.

**Exclude IPv6 addresses:** Specify IPv6 addresses of AAAA records that the appliance treats as nonexistent. The DNS server does not return the AAAA record of an address from this list. Instead, it synthesizes an AAAA record from the A record.

- **None:** Select this if you do not want to define specific IPv6 addresses or networks of AAAA records that the appliance treats as nonexistent. The appliance treats all IPv6 addresses as nonexistent. This is selected by default.
- **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance synthesizes

AAAA records from A records for the clients that have the **Allow** permission in the list. You can click **Clear** to remove the selected named ACL.

- **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
    - **IPv6 Address:** Select this to add an IPv6 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
    - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
      - **Permission:** Select **Allow** or **Deny** from the drop-down list.
    - **Any Address/Network:** Select this to allow or deny any IP addresses of AAAA records that the appliance treats as nonexistent.  
After you have added access control entries, you can do the following:
      - Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
      - Reorder the list of ACEs using the up and down arrows next to the table.
      - Select an ACE and click the Edit icon to modify the entry.
      - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
  - **Extensible Attributes:** You can modify the attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions:** This tab displays if you logged in as a superuser. For information, see [About Administrative Permissions](#).
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Enabling DNS64 Service

You can enable DNS64 at the Grid, member, and DNS view level. At least one DNS64 synthesis group must be configured before you can enable DNS64.

To enable DNS64 and apply DNS64 synthesis groups:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox -> Edit icon.  
**DNS View:** From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *dns\_view* checkbox -> Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the *Grid and Member DNS Properties* editor, click **Toggle Advanced Mode**, and then click **DNS64**. In the *View DNS Properties* editor, just click **DNS64**.
3. Do the following in the **DNS64** tab:
  - **Enable DNS64:** Select this checkbox.
  - **Synthesis Groups:** Click the Add icon and select a synthesis group.
4. Save the configuration and click **Restart** if it appears.

## DNS Record Scavenging

The DNS scavenging feature allows you to remove unused DNS resource records from zone data to prevent the accumulation of unneeded records. A scavenging operation determines, based on predefined rules, which records are not needed, i.e. are reclaimable, and removes them. For information about scavenging rules, see [Scavenging Rules](#) in this topic..

Scavenging is used for records with the dynamic record source type. Dynamic records are those created automatically, for example, via a dynamic DNS update. Static records, i.e. records that you add manually, can be identified as reclaimable based on scavenging rules but are not subject to scavenging. You can see the source type for each record in the DNS Resource Records viewer in Grid Manager.

You can use the records scavenging feature at the following levels in NIOS:

- Grid: scavenging is performed in all views and zones of the Grid.
- DNS view: scavenging is performed in all zones of the view.
- Authoritative zone (a Grid primary or unassigned zone): scavenging is performed in the specified zone, but not in the subzones.

You can either scavenge DNS records immediately or schedule automatic scavenging. For more information, see [Scavenging DNS Records Immediately](#) and [Scheduling Automatic Scavenging](#) in this topic..

You can organize and monitor records identified as reclaimable by using Smart Folders. For information, see [Smart Folders](#).

Scavenging events are logged in the NIOS syslog. You can view it, as described in [Viewing the Syslog](#) and [Searching in the Syslog](#).

The records are removed to the Recycle Bin and can be restored from there. For more information, see [Restoring Reclaimed Records](#) in this topic..



#### Note

Membership in the DNS Admin group is required to complete scavenging operations. For details, see [Administrative Permissions for DNS Records Scavenging](#) below. See [Forcing Creation Timestamp Initialization for Unchanged Records](#) for information on handling the creation timestamp of records that remain unchanged at DDNS updates.

## Scavenging Rules

You can configure the following match rules to identify reclaimable DNS resource records:

- **Resource Record Type:** This rule allows you to specify the record type to run scavenging on. A record is reclaimable if its type matches or does not match the type specified in the rule. The record types that support scavenging include the following:
  - A
  - AAAA
  - PTR
  - CNAME
  - DNAME
  - MX
  - SRV
  - NAPTR
  - TXT
- **Creation Time:** This rule allows you to identify reclaimable records based on the record's creation timestamp. You can see the "Creation Time" value for the records in the DNS Resource Records viewer. Note that in the case of an upgrade to NIOS 7.3, the creation time is not initialized. Therefore, the "Creation Time" rule does not apply to the records created before the upgrade.
- **Last Queried Time:** This rule allows you to identify reclaimable records based on when they were last queried for their DNS data. You can see the last queried time for the records in the DNS Resource Records viewer. Note that if you use this rule, also select **Enable last queried time monitoring for resource records** in the Grid, view, or zone scavenging properties, as described in the next section.
- **Last Discovered Time:** This rule allows you to identify reclaimable records based on the record's last discovered timestamp. This rule is applicable to A, AAAA, and PTR records.
- **Record Source:** This rule allows you to specify the record source – static or dynamic – to be used as a filter when identifying reclaimable records.
- **Associated Records:** This rule allows you to identify reclaimable records based on whether they have or do not have associated records. Record associations are supported for address records (A, AAAA, and PTR). Additionally, you can reclaim the associated records when reclaiming the original ones by enabling the option

**When reclaiming A, AAAA, or PTR records, also reclaim the corresponding, symmetric A, AAAA, and PTR records** in the scavenging properties, as described in the next section.

- **Extensible Attributes:** You can specify extensible attributes that reclaimable records should match in addition to the scavenging rules described above.

## Configuring DNS Record Scavenging Properties

You can configure the DNS record scavenging properties at the Grid, DNS view, or DNS zone level. According to the NIOS inheritance pattern for object properties, the scavenging properties configured at a given level are inherited by the level below, unless overridden.

To configure the DNS record scavenging properties, complete the following:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**DNS view:** From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> dns\_view checkbox -> Edit icon.  
**DNS zone:** From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> click a DNS view -> zone checkbox -> Edit icon.
2. If the properties editor is in basic mode, click **Toggle Advanced Mode**.
3. Click **DNS Scavenging**.
4. **Enable last queried time monitoring for resource records:** Select this if you are going to use the Last Queried Time rule. This setting enables monitoring the time when the resource record was last queried for its DNS data. For more information on DNS queries monitoring for resource records, see [Monitoring DNS Queries](#).
5. **Enable last queried time monitoring for zones:** This setting enables monitoring the time when the zone, or at least a single record in it, was last queried for its DNS data. The data resulting from zone last queries time monitoring is displayed in the zones viewer (**Data Management** -> **DNS** -> **Zones** -> click a *DNS view* to open zones list). Note that enabling monitoring for a zone does not enable monitoring for child zones.
6. You can configure the set of ACLs (Access Control Lists) to filter clients on DNS queries from updating the last-queried timestamp, under **Prevent the following ACLs or ACEs from updating the last queried timestamp**. To configure the ACLs, you should select either **Enable last queried time monitoring for resource records** or **Enable last queried time monitoring for zones** option, these options are disabled by default.
7. Select one of the following:
  - **None:** Select this option if you do not want to configure any access control for updating the last queried time stamp. When you select this option, NIOS will allow updates to the last queried time stamp for the queries received from any client. This is selected by default.
  - **Named ACL:** Select this option and click **Select Named ACL** to select a named ACL that you want to use. If you have only one named ACL, only that named ACL is displayed. When you select this option, the appliance prevents clients with the **Include** permission from updating the last queried timestamp. You can click **Clear** to remove the selected named ACL.
  - **Set of ACEs:** Select this option to configure individual access control entries (ACEs). Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so that you can specify additional information about the item you are adding.
    - **IPv4 Address and IPv6 Address:** Select this option to add an IPv4 or IPv6 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Include** by default. You can change it to **Exclude** by clicking the field and selecting **Exclude** from the drop-down list. When you select **Include**, the appliance prevents the client from updating the last queried timestamp. When you select **Exclude**, the appliance allows the client to update the last queried timestamp.
    - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
      - **Permission:** Select **Include** or **Exclude** from the drop-down list.
    - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
      - **Permission:** Select **Include** or **Exclude** from the drop-down list.
    - **Any Address/Network:** Select this option to include or exclude all IP addresses and networks to the last queried ACL list. The default permission is **Include**, which means the appliance prevents



updating the last queried timestamp from all clients. You can change this to **Exclude** to allow all clients to update the last queried timestamp.

- After you have added access control entries, you can perform the following:
  - Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
  - Reorder the list of ACEs using the up and down arrows next to the table.
  - Select an ACE and click the Edit icon to modify the entry.
  - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
- 8. Select **Enable record scavenging**.
- 9. To override the inherited properties, click **Override** and complete the fields.
- 10. Under **Match the following rule**, create a rule as follows. For information about rules, see *Scavenging Rules* above.
  - **Choose Filter:** Select a criterion from the drop-down list.
  - **Choose Operator:** Select an operator for the filter criterion.
  - In the value field, enter the value for the filter field. To add another rule:
  - Click **+** to add another rule at the same level.
  - Click **|<-** to add an **all** (logical AND) or **any** (logical OR) operator line and a parenthetical rule that is indented one level and above the first rule.
  - Click **->|** to add an **all** (logical AND) or **any** (logical OR) operator line and a parenthetical rule that is indented one level.  
To logically combine the whole ruleset, select **Match all of the following rules** or **Match any of the following rules**.  
After you add all the match rules, you can click **Reset** to remove the previously configured rules and start again.
- 11. Under **Match records with the following extensible attribute**, add an extensible attribute to use as an additional criterion for finding necessary records.
  - **Choose Operator:** Select an operator for the filter criterion.
  - **Choose Filter:** Select a criterion from the drop-down list.
  - In the value field, enter the value for the filter field.

To add another extensible attribute, click **+**.

1. To logically combine the extensible attributes set, select **Match all records with the following extensible attributes** or **Match any records with the following extensible attributes**.
2. After you add all the extensible attributes, you can click **Reset** to remove the previously configured attributes and start again.  
Note that the extensible attributes rule is always combined with the rest of the match rules using the AND operator.
3. **When reclaiming A, AAAA or PTR records, also reclaim the corresponding, symmetric A, AAAA and PTR records:** Select this if you want to reclaim records associated to the ones identified as reclaimable.
4. To configure a schedule for automatic records scavenging, select **Enable scheduled record scavenging**. See *Scheduling Automatic Scavenging* in this topic.
5. Click **Save & Close** or **Save**.

## Scheduling Automatic Scavenging

You can schedule a scavenging operation only at the Grid level. For a scavenging operation at the view or zone level, you can use the schedule inherited from the Grid.



### Note

Infoblox recommends manually testing the configured scavenging settings before enabling scheduled scavenging.

1. In the DNS record scavenging properties described in the previous section, select the **Enable scheduled record scavenging** checkbox.
2. To enable automatic scavenging after records are marked as reclaimable, select **After marking a record as reclaimable, automatically reclaim the record**.
3. Click the Scheduling icon and complete the following in the Scavenging Scheduler dialog:
  - Select how often you want to execute the scavenging. You can select **Once, Hourly, Daily, Weekly, or Monthly**.
  - If you select **Once**, complete the following:
    - Enter the day in the date picker and select a month from the drop-down list.
    - Enter a time in the hh:mm:ss AM/PM format. You can also select a time from the drop-down list.
    - Choose the **Time Zone**.
  - If you select **Hourly**, complete the following:
    - **Schedule every hour(s) at**: Enter the number of hours between each scavenging instance. You can enter a value from 1 to 24.
    - **Minutes past the hour**: Enter the number of minutes past the hour. For example, enter 5 if you want to schedule the scavenging operation five minutes after the hour.
    - Choose the **Time Zone**.
  - If you select **Daily**, complete the following:
    - Click either **Every day** or **Every weekday**.
    - Enter a time in the hh:mm:ss AM/PM format. You can also select a time from the drop-down list.
    - Choose the **Time Zone**.
  - If you select **Weekly**, complete the following:
    - **Schedule every week on**: Select any day of the week.
    - Enter a time in the hh:mm:ss AM/PM format. You can also select a time from the drop-down list.
    - Choose the **Time Zone**.
  - If you select **Monthly**, complete the following:
    - **Schedule the day of the month**: Enter the day of the month and the monthly interval. For example, to schedule the rule update on the first day after every 2 months, you can enter Day 1 every 2 month(s).
    - Enter a time in the hh:mm:ss AM/PM format. You can also select a time from the drop-down list.
    - Choose the **Time Zone**.
4. Click **OK**.

## Scavenging DNS Records Immediately

To perform record scavenging for the Grid, a DNS view, or a zone according to the predefined rules, use the Scavenge Records command from the Toolbar. This adds a background task that starts immediately or, if another scavenging task is in progress, after its completion.

The scavenging is split into two stages that you can execute separately or together:

- **Mark records as reclaimable**: This stage analyzes the records against the scavenging rules. The records matching the rules are marked as reclaimable, i.e. their "Reclaimable" flag is set to "Yes" in the DNS Resource Records viewer. These records can be reclaimed by using the second stage, unless you disable scavenging for them as described in *Disabling Scavenging for Individual Records* in this topic.
- **Reclaim records marked as reclaimable**: This stage automatically removes the records marked as reclaimable in the result of the execution of the first option. Running only the "Reclaim records marked as reclaimable" stage without the analysis stage does not perform a new analysis on the affected records. It only removes the records marked as reclaimable during the previous analysis.

Also, you can reset the reclaimable flag of the records. As an example of when this may be useful: if records have previously been marked as reclaimable and under a revised scavenging policy some records may no longer be reclaimable.



### Note

To start immediate scavenging of DNS records, you must first carefully define the scavenging properties, as described in *Configuring DNS Record Scavenging Properties* in this topic.



To scavenge DNS records immediately, complete the following:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Scavenge Records** -> **Scavenge Grid Records**.  
**DNS view:** From the **Data Management** tab, select the **DNS** tab, click a DNS view, expand the Toolbar, and then click **Scavenge Records** -> **Scavenge View Records**.  
**DNS zone:** From the **Data Management** tab, select the **DNS** tab, click a DNS view, click a zone, expand the Toolbar, and then click **Scavenge Records** -> **Scavenge Zone Records**.
2. Select one of the following:
  - **Scavenge Records:** Select this to proceed to the record scavenging. Go to the next step.
  - **Reset reclaimable flag:** Select this to set the "Reclaimable" flag of all affected records to "No".
3. If you chose **Scavenge Records**, select one of the following options or both of them:
  - **Mark records are reclaimable**
  - **Reclaim records marked as reclaimable**Note that static records are never reclaimed automatically even if they are marked as reclaimable. You can only delete static records manually from the DNS Resource Records viewer.
4. Click **Start**.

To check the progress of the current scavenging task, you can use the DNS Record Scavenging widget in the Dashboard. You can also view a scavenging report, as described in [DNS Scavenged Object Count Trend](#).

The scavenging task may be subject to an approval workflow. For information on approval workflows, see [Configuring Approval Workflows](#).



#### Note

Keep in mind that the **Enable record scavenging** property for a lower scavenging scope (e.g. view or zone) can override this property for the upper scope (i.e., Grid or view respectively). For example, if you run scavenging on the Grid with the scavenging option disabled, and there are some views or zones on which scavenging is enabled, this results in the records of the affected views and zones being scavenged. Vice versa, if scavenging is disabled for certain views or zones and you run scavenging on the Grid with the scavenging option enabled, the corresponding views and zones are excluded from scavenging.

## Disabling Scavenging for Individual Records

You can disable scavenging for individual records, even if they are marked as reclaimable. In this case, the record is never reclaimed unless you enable the scavenging for it again.

To disable scavenging for a record, complete the following:

1. In the *DNS Resource Records* viewer, select the appropriate record.
2. Click **Edit**.
3. In the record properties dialog, click **DNS Scavenging**.
4. Select the **Disable scavenging for this record** checkbox.
5. Click **Save & Close**.  
Additionally, you can see the following information in the resource record scavenging properties:
  - Record creation time
  - Record last queried time
  - Whether the record is reclaimable
1. For records synced from MS servers, the creation timestamp is not synced. This implies the following limitations:
  - When a zone is converted from MS to NIOS, the timestamp is initialized to the time when the operation occurs.
  - When a zone is converted from NIOS to MS, the timestamp is reset.

## Administrative Permissions for DNS Records Scavenging

By default, only superusers can perform DNS records scavenging. Limited-access users can use the scavenging functionality if they have the corresponding DNS scavenging permissions. For more information about admin permissions, see [About Administrative Permissions](#).

The DNS scavenging permissions are global to Grid Manager. They are used in addition to the regular DNS global and object permissions. For more information about the DNS permissions, see [Administrative Permissions for DNS Resources](#).

The following operations require scavenging permissions:

- Modifying scavenging properties for the Grid, a view, or a zone
- Configuring a scavenging schedule
- Performing a scavenging task
- Viewing the DNS Record Scavenging dashboard widget
- Viewing the DNS Scavenged Object Count Trend report

## Restoring Reclaimed Records

A reclaimed record remains in the Recycle Bin until the bin is emptied. You can restore the deleted records from the Recycle Bin, as described in [Restoring Objects from the Recycle Bin](#).

The Recycle Bin does not display information on whether a record was deleted during a scavenging process or manually. Therefore, you cannot restore the reclaimed data only.

When a record is restored from the Recycle Bin, its Reclaimable flag is reset to "No".



### Note

Only a super user can restore records reclaimed during a recurring scavenging task.

## Monitoring DNS Queries

You can monitor DNS resource records by their **Last Queried** time. You can configure this feature in the *Grid DNS Properties* editor -> **DNS Scavenging** tab. Infoblox recommends that you keep the number of zones or domains for monitoring below 1000; specifying more may adversely affect performance.

To view DNS queries by their **Last Queried** time:

1. From the **Data Management** tab -> **DNS** tab, click the **Query Monitoring** tab.
2. The **Query Monitoring** tab provides a **Last Queried** report for the monitored resource records, including the following information:
  - **Network View**: Network view name. You cannot sort on this column. This column is hidden by default.
  - **DNS View**: DNS view name. This column is hidden by default. **Zone**: FQDN of zone.
  - **Name**: FQDN of resource record.
  - **Record Type**: Resource record type.
  - **Record Data**: Value of resource record, such as address of an A record.
  - **Monitored Since**: Date monitoring started.
  - **Last Queried**: Displays "Not Monitored", "Not Queried Since xxxx", or date of last query.  
Note this report does not display the last queried information for automatically generated NS records.

To enable last queried time monitoring for resource records, do the following in the DNS scavenging properties for the Grid, a view, or a zone:

1. **Grid**: From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**DNS view**: From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> `dns_view` checkbox -> Edit icon.

**DNS zone:** From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> click a DNS view -> *zone* checkbox -> Edit icon.

2. If the properties editor is in basic mode, click **Toggle Advanced Mode**.
3. Click **DNS Scavenging**.
4. Select the **Enable last queried time monitoring for resource records** checkbox. For more information, see [Configuring DNS Record Scavenging Properties](#).
5. Click **Save & Close**.



#### Note

Exporting the query monitoring data may take longer than usual if the report contains a lot of records. Also, if a Grid secondary server uses zone transfer to update zone data from a Grid primary server, NIOS does not monitor queries made to the Grid secondary server and it does not update the last queried timestamp for the resource records in a zone.

When multiple values are specified with the same filter, the filter applies *or* logic, e.g. 'a' or 'b'. Other perspectives in NIOS UI apply *and* logic, e.g., 'a' and 'b'. You can use the following filters to get specific information in this report:

- **DNS View:** DNS view name.
- **Not Queried:** Specify a date when the last query was made. The only operator is "Since".
- **Zone:** FQDN of zone.
- **Type:** Only a single record type filter can be specified. This filter has the following resource records:
  - A Records
  - AAAA Records
  - BulkHost
  - CAA Records
  - CNAME Records
  - DNAME Records
  - DS Records
  - DTC LBDN Records
  - Host Address
  - Host Alias
  - Host Record
  - MX Records
  - NAPTR Records
  - NS Records
  - PTR Records
  - Resource Record
  - SRV Records
  - Shared A Record
  - Shared AAAA Record
  - Shared CNAME Record
  - Shared MX Record
  - Shared SRV Record
  - TXT Records
  - Other Records



#### Note

NIOS does not monitor queries or update timestamp for DNSSEC records, except for DS records. As a result, the **QueryMonitoring** tab displays "Not Monitored" in the Last Queried column for all DNSSEC records. In addition, the Not Queried filter does not display any DNSSEC records.

## Configuring DNS Traffic Control Properties

You can configure the DNS Traffic Control properties at the Grid or member level. The member DTC properties are inherited from the Grid DTC properties unless you override them.

The following sections explain how to configure DTC properties for the Grid or a Grid member.

### Configuring Grid DNS Traffic Control Properties

To configure DNS Traffic Control properties for the Grid, complete the following:

1. From the **Data Management** tab, select the **DNS** tab, expand the toolbar, and then click **Grid DNS Properties**.
2. In the **Traffic Control** tab, complete the following:
  - **Extensible Attributes Source Types for Topology Rules:** Specify up to four extensible attributes to use as source types when defining DNS traffic control topology rules. For information about the extensible attribute topology rules, see [Defining Topology Rulesets](#). You can use either predefined EAs or use your own. When using the predefined EAs you can select all or select specific types of IPAM objects, network containers, networks, ranges, and hosts and their EA values to be considered for the predefined EAs. For information about predefined extensible attributes, see [. For information about creating custom extensible attributes, see Adding Extensible Attributes.](#)  
You can select only string, enum, and integer types as the EAs to be used. Also, you cannot select EAs for which multiple values have been set. That is, the **Allow multiple values** checkbox must not be enabled for EAs that you want to use to define topology rules.  
Note that the enabled IPAM object for Extensible Attributes Source Types for Topology Rules depend on the upgrade history of the Grid. In case, Grid is upgraded from the earlier NIOS version to a later version which supports different IPAM object types, only IPAM Networks and IPAM Ranges types are enabled by default. However, in a Grid that is created with the possibility to select different IPAM object types, all the IPAM object types are enabled by default.
  - **When DNS Traffic Control is enabled, direct traffic according to EDNS0 Client Subnet when possible:** Select this checkbox to direct traffic according to EDNS0 client subnet option when DNS Traffic Control is processing DNS queries.  
You can enable the appliance to redirect traffic according to EDNS0 client subnet option when DNS Traffic Control is processing DNS queries that contain the EDNS0 client subnet option. When you enable this feature, DNS Traffic Control querying process uses the client address specified in the EDNS0 client subnet option of the DNS query and the appliance includes the EDNS0 client subnet option in the response message. If there are multiple EDNS0 client subnet options in a query, the appliance considers only the first option and ignores the other options. When this feature is disabled, DNS Traffic Control querying process ignores the EDNS0 client subnet option. For more information about EDNS0, see [Using Extension Mechanisms for DNS \(EDNS0\)](#).
  - **Specific behavior for DNS queries:** Select one of the following options if you want the appliance to return DNS responses when no DNS traffic control responses are available. The **Return DNS response if there are no DNS Traffic Control responses available** option is selected by default.
    - **Return DNS response if there are no DNS Traffic Control responses available:** Select this option if you want the appliance to return DNS responses when DNS Traffic Control responses are not available.
    - **Drop LBDN matched DNS queries during full health update:** Select this option to drop all LBDN queries when the DNS service is waiting to receive a full health status update from the health monitor. The appliance drops the LBDN queries and returns a SERVFAIL response.
    - **No specific behavior:** Select this option when you do not want the appliance to return DNS responses when DNS Traffic Control responses are not available.
  - **Return the following type of response from DNSSEC signed zones:** Select one of the following response types for DNSSEC-signed zones:
    - Signed
    - Unsigned

For more information on the Signed and Unsigned modes, see [Managing LBDN Records](#).
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring Member DNS Traffic Control Properties

To configure DNS Traffic Control properties for a Grid member, complete the following:

1. From the **Data Management** tab, click the **DNS** tab -> **Members** tab -> *member* checkbox -> Edit icon.  
Alternatively, you can click the **DNS** tab -> **Traffic Control** tab, click **Member DNS Properties** in the toolbar, and then select the required member.
2. In the **Traffic Control** tab, complete the following:
  - **DTC Health Check Source:** Select which NIOS network interface to use for the DTC health check. Choose one of the following:  
Note that for vNIOS appliances, some of the options may vary depending on your vNIOS configuration. For example, if you are using a single network interface instance of vNIOS for GCP, you will see choices specific to the LAN1 interface only. For more information, see the vNIOS documentation specific to your product at [Appliances](#).
    - ANY interface
    - VIP interface
    - LAN2
    - MGMT interface
    - IP (This is displayed only when you have configured additional IP addresses in the network settings. Specify the IP address of the source.)
  - **Specific Behavior for DNS queries:** Select one of the following options if you want the appliance to return DNS responses when no DNS traffic control responses are available. The **Return DNS response if there are no DNS Traffic Control responses available** option is selected by default. To override the Grid setting, click **Override**.
    - **Return DNS response if there are no DNS Traffic Control responses available**
    - **Drop LBDN matched DNS queries during full health update**
    - **No specific behavior**For more information, see [Configuring Grid DNS Traffic Control Properties](#) above.
  - **When DNS Traffic Control is enabled, direct traffic according to EDNS0 Client Subnet when possible:** To retain the same setting as the Grid, keep the inherited value. To override the Grid setting, click **Override**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## DNS Views

DNS views enable the NIOS appliance to serve different versions of DNS data based on the host accessing it. The topics in this section include:

- [About DNS Views](#)
- [Configuring DNS Views](#)
- [Configuration Example: Configuring a DNS View](#)

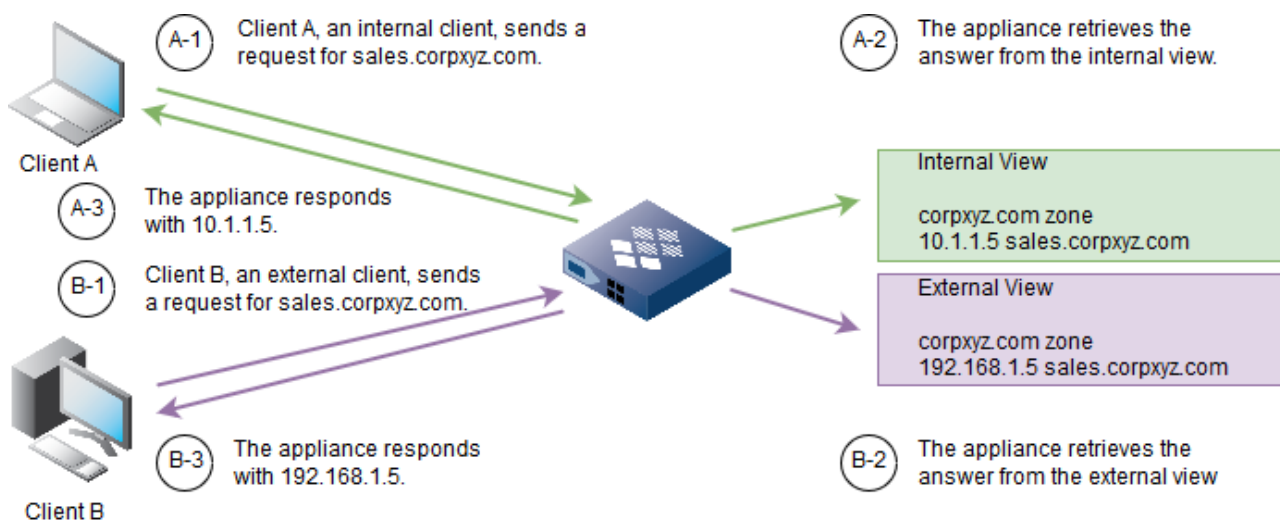
## About DNS Views

DNS views enable the NIOS appliance to serve different versions of DNS data based on the host accessing it.

DNS views provide the ability to serve one version of DNS data to one set of clients and another version to another set of clients. With DNS views, the NIOS appliance can provide a different answer to the same DNS query, depending on the source of the query.

In the below figure, the appliance has two views: an Internal and an External DNS view. When the appliance receives queries from DNS clients, it responds with data from either the Internal or External DNS view, depending on the source IP address. When the appliance receives a query from Client A and determines that it can resolve the query from data in the Internal view, the appliance responds with the IP address of the site in the Internal view. When the appliance receives a query from Client B and determines that it can resolve the query from data in the External view, it responds with the IP address in the External view.

*Internal and External Views*



You can configure both forward and reverse mapping zones in DNS views and provide DNS services, such as name resolution, zone transfers and dynamic DNS updates. For information about these services, see [Configuring DNS Services](#).

You can provide multiple views of a given zone with a different set of records in each DNS view. In the below figure, both views contain the corpxyz.com zone and the sales.corpxyz.com zone. The finance.corpxyz.com zone is only in the internal DNS view, and only internal users are allowed to access records in that zone. Resource records can also exist in multiple zones. In the example, the A records for serv1.sales.corpxyz.com and serv2.sales.corpxyz.com are in the sales.corpxyz.com zones in both views.

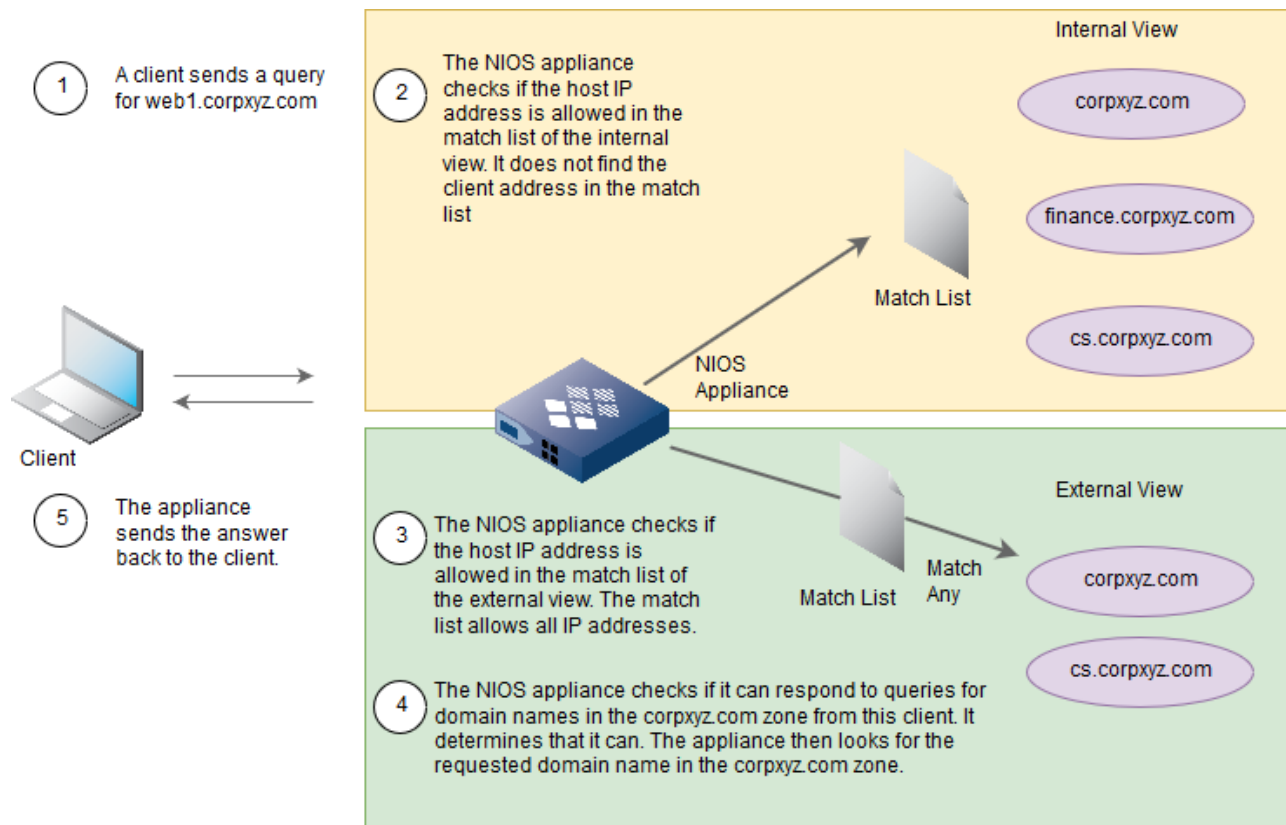
Zone Data in Each DNSView

Internal DNS View	corpxyz.com		sales.corpxyz.com		finance.corpxyz.com
MX	rmail.corpxyz.com	A	serv1.sales.corpxyz.com	A	server.finance.corpxyz.com
NS	dnsoneA.corpxyz.com	A	serv2.sales.corpxyz.com	A	printer.finance.corpxyz.com
A	host1.corpxyz.com	A	serv3.sales.corpxyz.com	A	fin1.finance.corpxyz.com
A	host2.corpxyz.com	AAA	printer.sales.corpxyz.com host1.sales.corpxyz.com host2.sales.corpxyz.com	A	fin2.finance.corpxyz.com
External DNS View	corpxyz.com		sales.corpxyz.com		
MX AA	email.corpxyz.com web1.corpxyz.com web2.corpxyz.com	AAA A	web3.sales.corpxyz.com ftp.sales.corpxyz.com serv1.sales.corpxyz.com serv2.sales.corpxyz.com		

You can control which clients access a DNS view through the use of a match list specifying IP addresses and/or TSIG (transaction signature) keys. When the NIOS appliance receives a request from a client, it tries to match the source IP address and/or TSIG key with its match list when determining which DNS view, if any, the client can access. After the appliance determines that a client can access a DNS view, it checks the zone level settings to determine if it can provide the service that the client is requesting. For information on TSIG keys or defining zone transfer settings, see [Enabling Zone Transfers](#). For more information on match lists, see [Defining Match Clients Lists](#). For information on defining query settings, refer to [Controlling DNS Queries](#).

The figure Query Resolution below illustrates how the NIOS appliance resolves a query for a domain name in a zone of a DNS view. In the example, the internal DNS view is listed before the external DNS view. Therefore, when the appliance receives a query, it checks the match list of the internal DNS view first. If it does not find the source address in the match list of the internal DNS view, it checks the match list of the external DNS view. The match list of the external DNS view allows all IP addresses. Next, the NIOS appliance checks the zone level settings to determine if it is allowed to resolve queries from the client for domain names in that zone. After the appliance determines it is allowed to respond to queries from this client, it resolves the query and sends back the response to the client.

### Query Resolution



When you create more than one DNS view, as shown in the figure Query Resolution above, the order of the views is important. View order determines the order in which the NIOS appliance checks the match lists. In the figure Query Resolution, the internal DNS view is listed before the external DNS view. If the views were reversed, no hosts would receive DNS replies from the internal DNS view because the match list of the external DNS view allows replies to clients with any IP address.

In a Grid, each Grid member can host its own set of views. A Grid member can serve as the primary or secondary server for multiple views of a particular zone. For information about specifying primary and secondary servers, see [Assigning Zone Authority to Name Servers](#).

### About DNS Views and Network Views

The NIOS appliance provides one default DNS view, which is always associated with the default network view. You can create additional network and DNS views. A network view is a single routing domain with its own networks. For information about network views, see [Managing IPv4 DHCP Data](#).

The default DNS view initially allows all IP addresses access, and has the same recursion setting as the Grid. You can change these properties and rename the default DNS view, but you cannot delete it. When you upgrade or migrate from a name server, or an earlier version of software that does not support DNS views, the appliance places all the zones defined in the older release in the default DNS view. You can then create additional views and organize the zones in each view.



When you create a network view, the appliance automatically creates a corresponding DNS view with "default." prepended to the name of the network view. You can rename the system-defined DNS view and configure its properties. If the appliance contains only one network view, all DNS views are associated with that network view. If there are 20 or less network views configured, the appliance displays the network views in the drop-down list on the left of the top navigation bar of the **Data Management** tab of Grid Manager. The appliance displays the *Network View Selector* dialog box if there are more than 20 network views configured. You can adjust the page size of the selector by choosing the number of network views to be displayed on each page from the **Page Size** drop-down list. If the number of network views exceeds the selected number, the selector displays the data on multiple pages. If you have a large number of network views, select a larger page size so you can quickly locate a network view without excessive paging through the list. The default page size is 10.

A DNS view can be in one network view only, but a network view can have multiple DNS views. If you enable dynamic DNS updates, you must specify which DNS view receives the updates. In a network view, only one DNS view can receive the dynamic DNS updates. For information, see [Sending DDNS Updates to a DNS Server](#).

## Configuring DNS Views

Following are the tasks to configure a DNS view:

1. Add a DNS view, as described in [Adding a DNS View](#) below.
2. Add zones to the DNS view. You can add authoritative forward-mapping and reverse-mapping zones, as well as delegated, forward, and stub zones. For information about configuring each type of zone, see [Configuring Authoritative Zones](#) and [Configuring Delegated, Forward, and Stub Zones](#).

You can optionally do the following:

1. Define a Match Clients list and a Match Destination list to restrict access to the DNS view. For more information, see [Defining Match Clients Lists](#) and [Defining a Match Destinations List](#) below.
2. Copy resource records from one zone to another. This is useful when different DNS views have the same zone and have multiple resource records in common. For information, see [Managing DNS Views](#) below.
3. Create resource records in a group and share the group among multiple zones. For information, see [Configuring Shared Record Groups](#).
4. Specify which interface IP address is published in the glue A record of the DNS view. For information, see [Changing the Interface IP Address](#).
5. Manage recursive views. For information, see [Managing Recursive DNS Views](#) below.
6. Manage the order of the DNS views, as this determines the order in which the NIOS appliance checks the Match Clients list. For information, see [Managing the Order of DNS Views](#).
7. Configure forwarders for a DNS view. For more information, see [Using Forwarders](#).
8. Enable AAAA filtering and configure a list of IPv4 networks and addresses for allowing or denying AAAA filtering from the appliance. For information, see [Controlling AAAA Records for IPv4 Clients](#).

### Adding a DNS View

You can add up to 1000 DNS views. When you add a DNS view, specify the following:

- The network view in which you are creating the DNS view.  
The appliance lists the network views only when there are multiple network views. Otherwise, it automatically associates the DNS view with the default network view.
- A Match Clients list specifying the hosts allowed access to the DNS view.  
If you do not define a list, the appliance allows all hosts to access the DNS view. For more information, see [Defining Match Clients Lists](#) below.
- Whether recursive queries are allowed.  
When a name server is authoritative for the zones in a DNS view, you can disable recursion since your name server should be able to respond to the queries without having to query other servers.  
if you want to allow a Grid member to respond to recursive queries from specific IP addresses, you can create an empty DNS view, that is, one that has no zones in it, and enable recursion. For information, see [Configuration Example: Configuring a DNS View](#).





#### Note

This setting overrides the recursion setting at the Grid and member levels..

To configure a new DNS view:

1. If there is more than one network view in the Grid, select the network view in which you are creating the DNS view.
2. From the **Data Management** tab -> **DNS** tab, expand the Toolbar and click **Add** -> **Add DNS View**.
3. In the *Add DNS View* wizard, complete the following fields:
  - **DNS View:** Enter the name of the DNS view. It can be up to 64 characters long and can contain any combination of printable characters. Each DNS view must have a unique name. You cannot create two DNS views with the same name, even if they are in different network views.
  - **Comment:** Optionally, enter information about the DNS view. You can enter up to 256 characters.
  - **Enable Recursion:** This field's initial default state is inherited from the Grid. It is inactive and greyed out until you click **Override**. After you click override, you can select or clear the checkbox to define a setting that applies to the DNS view only.

Note that a DNS view actually inherits its recursion setting from the Grid members that serve its zones. When you first create a DNS view though, it does not have any zones and therefore inherits its setting from the Grid. After you create zones in the DNS view, Grid Manager can then determine the associated members and display the resulting inheritance. If a DNS view has multiple zones served by multiple members with different recursion settings, you can view the different settings in the Multi-Inheritance viewer.

You can click **Inherit** to have the DNS view inherit its recursion setting from the Grid.
4. Save the configuration and click **Restart** if it appears at the top of the screen, or click **Next** to define a Match Clients list. For information, see [Defining Match Clients Lists](#) below or Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

### Defining Match Clients Lists

When you configure a DNS view, you can create a Match Clients list to identify source IP addresses and TSIG keys that are allowed or denied access to the DNS view. The NIOS appliance determines which hosts can access a DNS view by matching the source IP address or TSIG key with its Match Clients list. After the appliance determines that a host can access a DNS view, it checks the zone level settings to determine whether it can provide the service that the host is requesting for that zone.

If you do not configure a Match Clients list, then all devices are allowed access to the DNS view. However, if you configure a Match Clients list, then only those devices in the list with "Allow" permission can access the DNS view. All other devices are denied access, including Grid members. Therefore, to allow a primary server of a zone to receive dynamic DNS updates from member DHCP servers, you must add the members to the Match Clients list as well. Note that if you "Deny" permission to certain IP addresses or networks, you must add the "Allow Any" permission at the end of the Match Clients list to ensure that all other IP addresses and networks that are not in the "Deny" list are allowed access to the DNS view. You can add individual ACEs (access control entries) or a named ACL (access control list) to the Match Clients list. For information about named ACLs and how to define them, see [Defining Named ACLs](#).

## Defining a Match Clients List for a DNS View

You can define a Match Clients list for a DNS view when you add a new DNS view (second step of the Wizard) or when you edit an existing DNS view.

To define a Match Clients list for an existing DNS view:

1. From the **Data Management** tab, click the **DNS** tab > **Zones** tab > *dns\_view* checkbox -> Edit icon. Or, if there is only one DNS view, for example the predefined default view, you can just click the Edit icon beside it.
2. In the *DNS View* editor, select the **Match Clients** tab, and select one of the following:
  - **None**: Select this if you want to configure a Match Clients list. The appliance allows all clients to access the DNS view. This is selected by default.
  - **Named ACL**: Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance allow access to the DNS view from sources that have the **Allow** permission in the named ACL. You can click **Clear** to remove the selected named ACL.
  - **Set of ACEs**: Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding.
    - **IPv4 Address** and **IPv6 Address**: Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
    - **IPv4 Network**: In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address**: Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
      - **Permission**: Select **Allow** or **Deny** from the drop-down list.
    - **IPv6 Network**: In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address**: Enter an IPv6 network address and select the netmask from the drop-down list.
      - **Permission**: Select **Allow** or **Deny** from the drop-down list.
    - **TSIG Key**: In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
      - **Key name**: Enter a meaningful name for the key, such as a zone name or the name of the client or Grid member. This name must match the name of the same TSIG key on other name servers.
      - **Key Algorithm**: Select either **HMAC-MD5** or **HMAC-SHA256**.
      - **Key Data**: To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.
    - **DNSone 2.x TSIG Key**: Select this when the other name server is a NIOS appliance running DNS One 2.x code. The appliance automatically populates the value of the key in the **Value** field. The **Permission** column displays **Allow** by default. You cannot change the default permission.
    - **Any Address/Network**: Select this to allow or deny any IP addresses to access the DNS view. After you have added access control entries, you can do the following:
      - Select the ACEs that you want to consolidate and put into a new named ACL. Click the **Create new named ACL** icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
      - Reorder the list of ACEs using the up and down arrows next to the table.
      - Select an ACE and click the Edit icon to modify the entry.
      - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
3. Save the configuration and click **Restart** if it appears at the top of the screen. You can also click the Schedule icon at the top of the editor to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone.

## Defining a Match Destinations List

You can define a Match Destinations list that identifies destination addresses and TSIG keys that are allowed access to a DNS view. When the NIOS appliance receives a DNS request from a client, it tries to match the destination address or TSIG key in the incoming message with its Match Destination list to determine which DNS view, if any, the client can access. After the appliance determines that a host can access a DNS view, it checks the zone level settings to determine whether it can provide the service that the host is requesting for that zone.

You can define a Match Destination list when you edit an existing DNS view as follows:

1. From the **Data Management** tab, click the **DNS** tab > **Zones** tab > dns\_view checkbox -> Edit icon. Or, if there is only one DNS view, for example the predefined default view, you can just click the Edit icon beside it.
2. In the *DNS View* editor, select the **Match Destinations** tab, and select one of the following:
  - **None**: Select this if you want to configure a Match Destinations list. The appliance allows all destination addresses to access the DNS view. This is selected by default.
  - **Named ACL**: Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance allows access to the DNS view from the destination addresses that have the **Allow** permission in the named ACL. You can click **Clear** to remove the selected named ACL.
  - **Set of ACEs**: Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
    - **IPv4 Address and IPv6 Address**: Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
    - **IPv4 Network**: In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address**: Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
      - **Permission**: Select **Allow** or **Deny** from the drop-down list.
    - **IPv6 Network**: In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address**: Enter an IPv6 network address and select the netmask from the drop-down list.
      - **Permission**: Select **Allow** or **Deny** from the drop-down list.
    - **TSIG Key**: In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
      - **Key name**: Enter a meaningful name for the key, such as a zone name or the name of the client or Grid member. This name must match the name of the same TSIG key on other name servers.
      - **Key Algorithm**: Select either **HMAC-MD5** or **HMAC-SHA256**.
      - **Key Data**: To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.
    - **DNSone 2.x TSIG Key**: Select this when the other name server is a NIOS appliance running DNS One 2.x code. The appliance automatically populates the value of the key in the **Value** field. The **Permission** column displays **Allow** by default. You cannot change the default permission.
    - **Any Address/Network**: Select this to allow or deny any IP addresses to access the DNS view.

After you have added access control entries, you can do the following:

  - Select the ACEs that you want to consolidate and put into a new named ACL. Click the **Create new named ACL** icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
  - Reorder the list of ACEs using the up and down arrows next to the table.
  - Select an ACE and click the Edit icon to modify the entry.
  - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.

3. Save the configuration and click **Restart** if it appears at the top of the screen. You can also click the Schedule icon at the top of the editor to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone.

### Enabling the Match Recursive Only Option

You can enable the match-recursive-only option for the DNS view. When you enable this option, only recursive queries from matching clients match the selected DNS view. This option can be used in conjunction with the match client list and match destination list. Ensure that you configure those options and the order of the DNS views accordingly if you want to also enable the match-recursive-only option.

To enable the match-recursive-only option, complete the following:

1. From the **Data Management** tab, click the **DNS** tab > **Zones** tab > dns\_view checkbox -> Edit icon. Or, if there is only one DNS view, for example the predefined default view, you can just click the Edit icon beside it.
2. In the *DNS View* editor, select the **General** tab -> **Advanced** tab, and select the following:
  - **Enable match recursive only option:** This option is disabled by default. Select this option to enable the match-recursive-only option for the DNS view. When you select this option, only recursive queries from matching clients match this view. Note that this option can be used in conjunction with the match-clients and match-destinations options. Ensure that you configure those options and the order of the DNS views accordingly if you want to also enable match-recursive-only.
3. Save the configuration.



#### Note

You can also enable or disable the match-recursive-only option for a specific DNS view on a specific member by using the CLI command **set enable\_match\_recursive\_only**. For information about this command, refer to the *Infoblox CLI Guide*.

### Copying Zone Records

Different views of the same zone may have a number of records in common. If this is the case, you can copy zone records between views and zones.



#### Note

You cannot copy shared records and records that the NIOS appliance automatically creates, such as NS records and glue A records.

To copy zone records between DNS zones and views:

1. From the **Data Management** tab -> **DNS** tab, click **Copy Records** from the Toolbar.
2. In the *Copy Records* dialog box, Grid Manager displays the last selected zone or the zone from which you are copying zone records in the **Source** field. Complete the following to copy records:
  - **Destination:** Click **Select Zone** to select the destination zone. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box from which you can select one. After you select the zone, Grid Manager displays the associated DNS view.
  - **Copy All records:** Select this option to copy all the zone records.
  - **Copy Specific Records:** Select this option to copy specific types of records. Select a resource record type from the Available column and click the right arrow to move it to the Selected column.
  - **Copy Options:** Select one of the following:
    - Delete all records in destination before copying the records: Select to delete all resource records in the destination zone before the records are copied.
    - Overwrite existing records: Select to overwrite existing resource records that have the same domain name owners as the records being copied.

### 3. Click **Copy & Close**.



#### Note

When you copy resource records between zones and there are pending scheduled tasks associated with these records, the appliance allows the copying of zone records before it executes the scheduled tasks.

## Managing the DNS Views of a Grid Member

A Grid member can serve zones in different DNS views. You can manage the DNS views associated with a Grid member as follows:

- You can specify which interface IP address is published in glue A records in the DNS view, as described in [Changing the Interface IP Address](#) below.
- You can assign an empty recursive view to a member, as described in [Managing Recursive DNS Views](#).
- You can control the list of DNS views as described in [Changing the Order of DNS Views](#).

### Changing the Interface IP Address

By default, a Grid member publishes its LAN address in glue A records in the DNS view. You can change this default for each DNS view associated with a member. You can specify the NAT IP address or another IP address. To specify the interface IP address for glue A records in a view:

1. From the **Data Management** tab, click the **DNS** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.
2. In the *Member DNS Configuration* editor, click **Toggle Expert Mode** if the editor is in basic mode, and then select the **DNS Views** tab.  
The *Address Of Member Used in DNS Views* table lists the default DNS view and DNS views with zones that are served by the member.
3. To change the address, click the entry in the Interface column of a DNS view, and select one of the following:
  - **NAT IP Address:** Select this to use the member NAT address for glue A records in a Grid setting. Select this when you want to notify the Grid Master that it should expect packets from this member on the NAT address, not the configured interface address. The Grid Master broadcasts this NAT address to all NAT members outside of its NAT group. Do not use this option for an independent appliance serving as a DNS server. Select **Other IP Address** to publish the NAT address for the independent appliance. For information about NAT compatibility, see [NAT Groups](#).
  - **Other IP Address:** Select this to specify another address for glue A records, or to publish the NAT address for an independent appliance. Enter the address in the **Address** field.  
Note that the 255.255.255.255 limited broadcast address is reserved. The appliance does not automatically create glue A records for this address. You can however create an NS record without the associated glue records.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

### Managing Recursive DNS Views

When you add a DNS view that has recursion enabled, the appliance resolves recursive queries from hosts on the Match Clients list of that view. If the DNS view contains zones and you delete those zones, the appliance retains the view in its configuration file, as long as recursion is enabled in the view. Such a view is called an empty recursive DNS view because it does not contain any zones. It enables the appliance to respond to recursive queries from the specified clients.

In a Grid, all members automatically store DNS views that have recursion enabled in their configuration files. If you do not want a Grid member to respond to recursive queries for clients in a particular DNS view, you can remove the view from the member's configuration file.

To delete or retain an empty recursive DNS view in the DNS configuration file of a Grid member:

1. From the **Data Management** tab, click the **DNS** tab > **Members** tab > *Grid\_member* checkbox -> Edit icon.

2. In the *Member DNS Configuration* editor, click **Toggle Expert Mode** if the editor is in basic mode, and then select the **DNS Views** tab.
3. The *Recursive Views Assigned to this Member* section lists the empty recursive DNS views. Move a DNS view to the *Selected* column to explicitly assign the view to the Grid member and include it in the DNS configuration file of the member. Move a DNS view to the *Available* column to remove it from the configuration file of the member. Empty recursive DNS views that you retain in the configuration file are automatically listed at the bottom of the list of DNS views. You can move them up on the list when you manually change the order of the DNS views, as described in *Managing the DNS Views of a Grid Member* below.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Managing the Order of DNS Views

When a member receives a query from a DNS client, it checks the Match Client lists in the order the DNS views are listed in the *Order of DNS Views* table of the **DNS Views** tab in the DNS Member editor. The NIOS appliance can order DNS views automatically, or you can order the DNS views manually. If you choose to have the appliance automatically update the order of the DNS views, it does so after each of the following events:

- Adding a DNS view to a member.
- Removing a DNS view from a member.
- Changing the address match list of a DNS view hosted by the member.

## About IP Addresses and the Order of DNS Views

NIOS appliances with both IPv4 and IPv6 enabled can contain both types of addresses in the Match Clients list. When you enable IPv6 on the appliance, the order of DNS views in the GUI may be affected. Views are ordered and sorted automatically based on Match Clients lists. Views with IPv6 enabled are sorted as follows:

- If the Match Clients lists of all views contain IPv4 addresses only—The appliance orders views based on IPv4 addresses.
- If the Match Clients lists of all views contain IPv6 addresses only—The appliance orders views based on IPv6 addresses.
- If the Match Clients list of one DNS view has IPv6 addresses and all other views have IPv4 addresses—The appliance orders views based on IPv4 addresses, and the IPv6 address is given lowest priority in the ordering.
- If the Match Clients list of one DNS view has IPv4 addresses and all other DNS views have IPv6 addresses—The appliance orders DNS views based on IPv6 addresses, and the IPv4 address is given lowest priority in the ordering.
- If the Match Clients list of one DNS view has both IPv4 and IPv6 addresses—The appliance orders DNS views based on both IPv4 and IPv6 addresses, but more priority is given to the IPv4 addresses in the ordering.

The DNS views are ordered based on the number of IP addresses that are matched by the Access Control Lists (ACLs). The order of the DNS view is as follows:

- ANY
- Large Network
- Small Network
- Multiple Addresses
- Single Address

The actual precedence of the order of the views is also based on the ACL elements:

- any match: precedence = `UINT_MAX + 1`
- address match: precedence += 1
- TSIG match: precedence += 1
- network match: precedence += 129 - split (BOTH v4 and v6)

Note that views with the same precedence are sorted based on the internal view name. For example, `'_default'` or `'0'`.



### Note

Only superusers can change the order of the views.

## Changing the Order of DNS Views

To change the order of DNS views:

1. From the **Data Management** tab, click the **DNS** tab > **Members** tab > *Grid\_member* checkbox -> Edit icon.
2. In the *Member DNS Configuration* editor, click **Toggle Expert Mode** if the editor is in basic mode, and then select the **DNS Views** tab.
3. In the Order of DNS Views section, select one of the following:
  - **Order DNS Views Automatically:** Click this to automatically order views after adding a new DNS view, removing a DNS view, or changing the match client list.
  - **Order DNS Views Manually:** This table lists the DNS views that have zones assigned to the Grid member and the empty recursive views associated with the member. Select a DNS view, then click an arrow to move it up or down in the list.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Managing DNS Views

You can list the DNS views, and then modify, disable, or remove any custom DNS view. You can modify and disable the default DNS view; however, under no circumstances can it be removed.

## Listing DNS Views

After you configure additional DNS views, you can list all DNS views by navigating to the **Data Management** tab -> **DNS** tab -> **Zones** panel. This panel lists DNS views only after you modify the default DNS view or add a DNS view. If you never added DNS views or modified the default DNS view, this panel does not display the default DNS view. Instead, it lists the zones in the default DNS view. To display the properties of the default DNS view and edit it, use the Global Search function to locate and edit it.

Note that if you have not used Grid Manager to add a new DNS view, and you import DNS views through the Data Import Wizard or the API, you must log out and log back in to Grid Manager to display the newly imported DNS views.

For each DNS view, this panel displays the following by default:

- **Comment:** Comments that were entered for the DNS view.
- **Site:** Values that were entered for this pre-defined attribute. You can also display the following column:
- **Disabled:** Indicates if the DNS view is enabled or disabled. Disabled DNS views are excluded from the named.conf file. You can double click a row and select the checkbox in this column to disable the network. Grid Manager displays a warning message when you select the checkbox. Click **Yes** to confirm or **No** to cancel. Note that disabling a DNS view may take a longer time to complete depending on the size of the data.

From this list, you can do the following:

- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).
- List the zones in a DNS view by clicking a DNS view name.
- Edit information about a DNS view, as described in the next section.
- Delete a DNS view, as described in [Deleting DNS Views](#) below.

## Modifying DNS Views

To modify a DNS view:

1. From the **Data Management** tab, click the **DNS** tab > **Zones** tab > *dns\_view* checkbox -> Edit icon.
2. In the *DNS View* editor, you can do the following:
  - In the **General** tab, you can change any of the information you entered through the wizard. You can also disable a DNS view to temporarily block access to a DNS view. Disabling a DNS view excludes it from the



- named.conf file. For a description of the fields, see the online Help or *Configuring DNS Views* Note that disabling a DNS view may take a longer time to complete depending on the size of the data.
- In the **Match Clients** tab, define or update a Match Clients list, as described in Defining Match Clients Lists.
  - In the Match Destinations tab, define or update match destinations, as described in Defining a Match Destinations List below.
  - In the **Forwarders** tab, configure forwards for the view, as described in [Using Forwarders](#).
  - In the **Queries** tab, enable AAAA filtering and configure a list of IPv4 networks and addresses for allowing or denying AAAA filtering, as described in [Enabling AAAA Filtering](#).
  - In the **DNSSEC** tab, you can specify parameters for DNSSEC as described in [Configuring DNSSEC on a Grid](#).
  - In the **Root Name Servers** tab, you can configure root name servers, as described in [About Root Name Servers](#).
  - In the **Sort List** tab, define a sort list for the DNS view, as described in [Defining a Sort List](#).
  - In the **Blacklist** tab, define blacklist rulesets, as described in [Enabling Blacklisting](#).
  - In the **Extensible Attributes** tab, you can modify the attributes. For information, see [Using Extensible Attributes](#).
  - The **Permissions** tab displays if you logged in as a superuser. For information, see [About Administrative Permissions](#).
  - In the **Record Scavenging** tab, define the rules for DNS records scavenging in the DNS view, as described in [Configuring DNS Record Scavenging Properties](#).
  - In the **Updates** tab, specify the secure dynamic updates settings, as described in [Secure Dynamic Updates](#).
3. Save the configuration and click **Restart** if it appears at the top of the screen.  
or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Deleting DNS Views

You can delete a DNS view if it is not the only view associated with a network view and if it is not selected for dynamic DNS updates. You cannot remove the system-defined default DNS view. When you remove a DNS view, the NIOS appliance removes the forward and reverse mappings of all the zones defined in the DNS view. To delete a DNS view:

- From the **Data Management** tab, select the > **DNS** tab> **Zones** tab-> *dns\_view* checkbox.  
To delete the DNS view immediately, click the Delete icon, and then click **Yes** to confirm the delete request. Grid Manager displays a warning message. Click **Yes** to continue or **No** to cancel the process. To schedule the deletion, click **Schedule Deletion** and in the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Deletions](#).  
Grid Manager moves the view to the Recycle Bin, from which you can restore or permanently delete it. Click **Restore** in the Recycle Bin to recover the deleted data. Click **Yes** in the Restore Item dialog box to restore or **No** to cancel the process. Note that deleting and restoring a DNS view may take a longer time to complete depending on the size of the data.

## Configuration Example: Configuring a DNS View

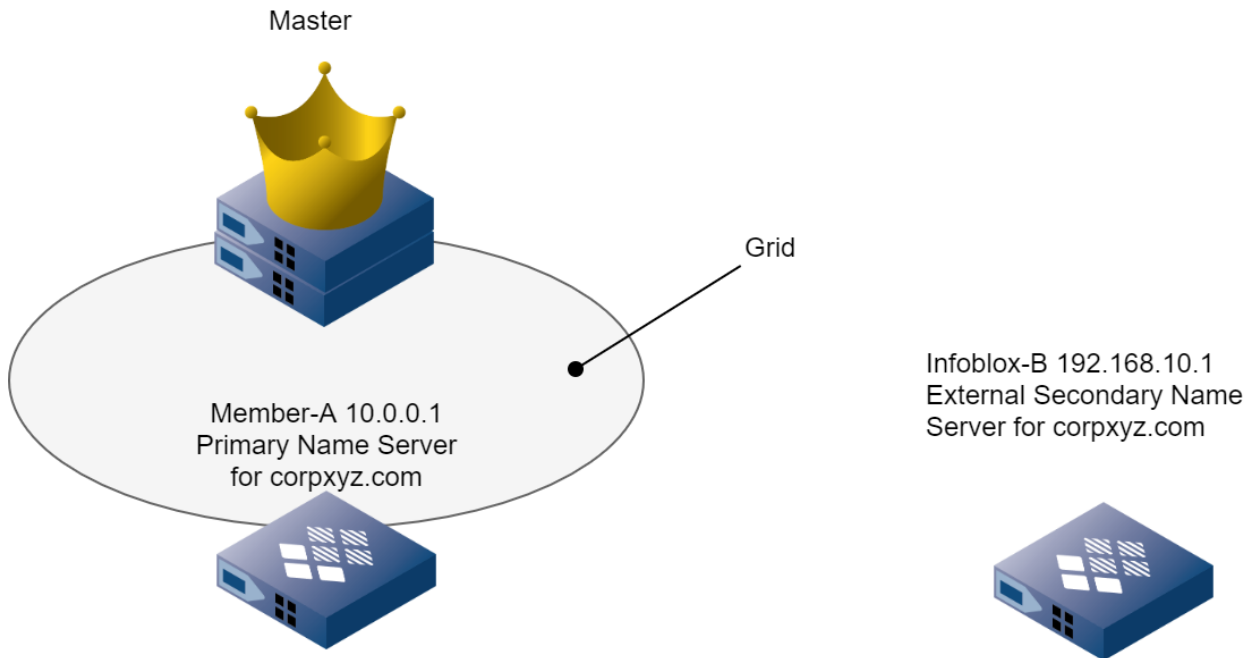
In the figure *Configuring a DNS View* below, Member-A is a member of a Grid. It is the primary name server for the corpxyz.com zone in the internal DNS view. It allows the IP address 192.168.10.1 and the 10.2.2.0/24 subnet access to DNS zone data in the internal DNS view. At the zone level, it allows transfers to an external secondary server, Infoblox-B, with an IP address of 192.168.10.1. Infoblox-B is a secondary server for the corpxyz.com zone. The process follows these steps:

1. [Adding an Internal DNS View on Member-A](#)
2. [Adding a Zone to a DNS View](#)



3. Copying Records Between DNS Zones from the corpxyz.com zone in the default DNS view to the corpxyz.com zone in the internal DNS view
4. Verifying the Configuration

#### Configuring a DNS View



#### Adding an Internal DNS View

1. Expand the Toolbar and click **Add** -> **Add DNS View**.
2. In the *Add DNS View* wizard, specify the following, and then click **Next**:
  - **Name:** internal
  - **Comment:** internal DNS view
3. In the *Match Clients* panel, click **Add** and select **IPv4 Network** from the drop-down list.
4. Do the following for IP addresses in the network 10.2.2.0/24:
  - Enter **10.2.2.0/24** in the in the **Address** field.
  - The **Permission** field displays **Allow** by default. Leave it as is.
  - Click **Add**.

You will have 255 allowed client addresses in the Match Clients list when you are done.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

#### Adding a Zone to a DNS View

1. Expand the Toolbar and click **Add** -> **Zone** -> **Add Auth Zone**.
2. In the *Add Auth Zone* wizard, click **Add an authoritative forward-mapping zone** and click **Next**.
3. Specify the following, and then click **Next**:
  - **Name:** Enter **corpxyz.com**.
  - **DNS View:** Select **Internal** from the drop-down list.
4. In step 3 of the wizard, do the following:
  - a. Select **Use this set of name servers**.
  - b. Click the Add icon and select **Grid Primary**.
  - c. Click **Select Member** and select **Member A** from the *Select Grid Member* dialog box.
  - d. Click **Add** to add the Grid member to the list of name servers.
  - e. Click the Add icon again and select **External Secondary**.
  - f. Enter the following information, and then click **Add**:
    - **Name:** Infoblox

- **IP Address:** 192.168.10.1
5. Click **Save & Edit** to display the *Authoritative Zone* editor and continue with the zone configuration.
  6. Click **Queries**.
  7. Click **Override**, and then click the Add icon and select **IPv4 Network**.
    - Enter **10.2.2.0/8** in the **Address** field.
    - The **Permission** field displays **Allow** by default. Leave it as is.
    - Click **Add**.
  8. This allows queries that the appliance answers from its internal DNS view.
  9. Save the configuration and click **Restart** if it appears at the top of the screen.

### Copying Records Between DNS Zones

1. Navigate to the default DNS view and select the corpxyz.com zone.
2. Expand the Toolbar and click **Copy Records**.
3. In the **Destination** field, click **Select Zone**, and then select the **corpxyz.com** zone in the Internal DNS view.
4. Select **Copy all records**, and then click **OK**.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

The records from corpxyz.com in the default DNS view are copied to corpxyz.com in the internal DNS view.



#### Note

Only superusers can copy A, AAAA, shared A, and shared AAAA records with a blank name. Limited-access users must have read/write permission to **Adding a blank A/AAAA record** in order to copy A, AAAA, shared A, and shared AAAA records with a blank name, otherwise the copying records operation might fail. You can assign global permission for specific admin groups and roles to allow to copy A, AAAA, shared A, and shared AAAA records with a blank name. For more information, see [Administrative Permissions for Adding Blank A or AAAA Records](#).

### Verifying the Configuration

1. In the **DNS** tab, click **Members** and select the **Member-A** checkbox .
2. Expand the Toolbar and click **View** -> **View DNS Configuration**.
3. In the *DNS Configuration File* viewer, scroll through the contents of the file.

Verify that the internal DNS view section is similar to the configuration file shown.

## Configuring DNS Zones

This section provides general information about DNS zones that you can configure and manage on the Infoblox appliance. The topics in this section include:

- [Configuring Authoritative Zones](#)
- [Creating an Authoritative Forward-Mapping Zone](#)
- [Creating an Authoritative Reverse-Mapping Zone](#)
- [Adding an Authoritative Subzone](#)
- [Locking, Unlocking, Enabling, Disabling Zones](#)
- [Configuring Authoritative Zone Properties](#)
- [Enabling Fixed RRset Ordering for NAPTR Records](#)
- [Configuring DNS Integrity Check for Authoritative Zones](#)
- [Domains and Zones](#)
- [Assigning Zone Authority to Name Servers](#)
- [Configuring Domain Controller List](#)
- [Importing Zone Data](#)
- [Enabling Zone Transfers](#)

- [Removing Zones](#)
- [Restoring Zone Data](#)
- [Configuring Delegated, Forward, and Stub Zones](#)
- [Viewing Zones](#)
- [Using Name Server Groups](#)
- [Viewing Name Server Groups](#)

## Configuring Authoritative Zones

An authoritative zone is a zone for which the local (primary or secondary) server references its own data when responding to queries. The local server is authoritative for the data in this zone and responds to queries for this data without referencing another server.

There are two types of authoritative zones:

- Forward-mapping – An authoritative forward-mapping zone is an area of domain name space for which one or more name servers have the responsibility to respond authoritatively to name-to-address queries.
- Reverse-mapping – A reverse-mapping zone is an area of network space for which one or more name servers have the responsibility to respond to address-to-name queries.

You can configure and manage authoritative forward-mapping and IPv4 and IPv6 reverse-mapping zones on an Infoblox appliance. In a Grid, an authoritative forward-mapping zone is an area of domain name space for which one or more Grid members have the responsibility to respond authoritatively to name-to-address queries. The Grid members can function as primary or secondary servers for the zone.

You can add arpa as the top-level forward-mapping zone and manage its resource records. You can also add in-addr.arpa (for ipv4 addresses) and ip6.arpa (for ipv6 addresses) as the top-level reverse-mapping zones.

You can create these top-level reverse-mapping zones under an arpa or a root parent forward-mapping zone or without a parent zone. If you want arpa, in-addr.arpa, and ip6.arpa zones on the appliance, you must manually create them. These zones are not auto-created.

Sample IPv4 reverse-mapping zone hierarchy:

```
. (root zone) > arpa > in-addr.arpa > 10.in-addr.arpa
```

Sample IPv6 reverse-mapping zone hierarchy:

```
. (root zone) > arpa > ip6.arpa > a.ip6.arpa
```

Following are the tasks to configure an authoritative zone:

1. Create the zone. The following sections explain how to create authoritative forward-mapping zones, reverse-mapping zones, subzones, and a custom root zone:
  - [Creating an Authoritative Forward-Mapping Zone](#)
  - [Creating an Authoritative Reverse-Mapping Zone](#)
  - [Creating a Root Zone](#)
2. Assign an Infoblox appliance as the primary or secondary server of the zone. For information, see [Assigning Zone Authority to Name Servers](#).
3. Import resource records or add resource records manually. The following provides information about resource records:
  - [Managing Resource Records](#)
  - [Importing Zone Data](#)
4. Configure additional parameters. For information, see [Configuring Authoritative Zone Properties](#).
5. Optionally, associate the zone with one or more networks. This is useful when you want to restrict the A, AAAA and host records to IP addresses from specific networks. For information, see [Associating Networks with Zones](#).

## Creating an Authoritative Forward-Mapping Zone

An authoritative forward-mapping zone is an area of domain name space for which one or more Grid members have the responsibility to respond authoritatively to name-to-address queries.



#### Note

A single forward-mapping zone can map names to both IPv4 and IPv6 addresses.

To create an authoritative forward-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar, and click **Add -> Zone -> Add Auth Zone**.
  2. In the *Add Authoritative Zone* wizard, click **Add an authoritative forward-mapping zone** and click **Next**.
  3. Specify the following:
    - **Name:** Enter the domain name for the zone. Omit the trailing period (" . ") that signifies the root zone. You can use IDNs as well. For information about IDNs, see [Support for Internationalized Domain Names](#).
    - **DNS View:** This field displays only when there is more than one DNS view in the current network view. Select a DNS view from the drop-down list.
    - **Comment:** Enter a descriptive comment about the zone.
    - **Disable:** Click this checkbox to temporarily disable this zone. For information, see [Enabling and Disabling Zones](#). Note that disabling a zone may take a longer time to complete depending on the size of the data.
    - **Lock:** Click this checkbox to lock the zone so that you can make changes to it and prevent others from making conflicting changes. For information, see [Locking and Unlocking Zones](#).
  4. Save the configuration, or click **Next** to continue to the next steps in the wizard as follows:
    - Define the name servers for the zone. For information on specifying primary and secondary servers, see [Assigning Zone Authority to Name Servers](#). For information on specifying authoritative name server groups, see [Using Authoritative Name Server Groups](#).
    - If you have assigned a Microsoft server as the primary server for the zone and if the zone is AD-integrated, you can configure a list of domain controllers that are allowed to add NS records to the zone. For information see, [Configuring Domain Controller List](#).
    - Define extensible attributes. For information, see [Using Extensible Attributes](#).
- or
- Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).
5. Click **Restart** if it appears at the top of the screen.

## Creating a Root Zone

The NIOS appliance allows you to create an internal root zone for your organization. When the appliance receives a query for DNS data that is not in its cache or authoritative data, it can query an internal root server after querying any specified forwarders. If you do not specify an internal root server and the appliance can access the Internet, it queries the Internet root servers. For information on root name server, see [About Root Name Servers](#).

To create a root zone, create an authoritative forward-mapping zone as described in [Creating an Authoritative Forward-Mapping Zone](#) above and specify the following:

- Enter a period (.) in the **Name** field.
- Optionally, enter a comment.
- Select a Grid member as the primary name server for the root zone.

Once created, the root zone automatically becomes the parent of all the zones under the root zone.

## Creating an Authoritative Reverse-Mapping Zone

An authoritative reverse-mapping zone is an area of network space for which one or more name servers—primary and secondary—have the responsibility to respond to address-to-name queries. Infoblox supports reverse-mapping zones for both IPv4 and IPv6 addresses. You can add in-addr.arpa and ip6.arpa as the top-level reverse-mapping zones. Note that you cannot add these zones using their IP addresses or netmasks, however, you can add them by name "in-addr.arpa" and "ip6.arpa" respectively.

## Specifying an RFC 2317 Prefix

*RFC 2317, Classless IN-ADDR.ARPA delegation* is an IETF (Internet Engineering Task Force) document that describes a method of delegating parts of the DNS IPv4 reverse-mapping tree that correspond to subnets smaller than a /24 (from a /25 to a /31). The DNS IPv4 reverse-mapping tree has nodes broken at octet boundaries of IP addresses, which correspond to the old classful network masks. So, IPv4 reverse-mapping zones (and delegation points) usually fall on /8, /16, or /24 boundaries.

With the proliferation of CIDR (Classless Inter-Domain Routing) support for routing, ISPs no longer assign entire /24 networks to customers that only need a handful of IPv4 addresses. In general, IPv4 address assignments no longer fall on classful boundaries. For DNS, a problem comes into play when an ISP gives a customer an address range that is smaller than a /24, but the customer also wants to be delegated the DNS reverse-mapping zone.

If the ISP gives you, for example, a subnet with a 25-bit mask, then you only have half of the /24 address range. If you configure your DNS server to be authoritative for the zone corresponding to a /24 subnet, the DNS server cannot resolve half of the possible reverse-mapping records in the zone. RFC 2317 defines an approach, considered a best practice, which addresses this issue.

In addition to IPv4 reverse-mapping zones, you can also configure IPv4 reverse-mapping delegation zones that have an RFC2317 prefix. For more information about configuring a delegation for a reverse-mapping zone, see [Configuring a Delegation](#).



### Note

Before enabling RFC 2317 support for zones, disable forwarders for the zone, especially when any sort of delegation (including RFC 2317) is being used. If you do not, reverse lookups may fail. For more information, contact Infoblox Support for the Tech Note on RFC 2317 delegation.

## Adding an IPv4 Reverse-Mapping Zone

To add an IPv4 reverse-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Auth Zone**.
2. In the *Add Authoritative Zone* wizard, click **Add an authoritative IPv4 reverse-mapping zone** and click **Next**.
3. Specify the following zone information:
  - Enter one of the following to identify the zone:
    - **IPv4Network:** Enter the IPv4 address for the address space for which you want to define the reverse-mapping zone and select a netmask from the **Netmask** drop-down list. Alternatively, you can specify the address in CIDR format, such as 192/8.  
To use an RFC 2317 prefix, select a netmask value that is between 25 to 31, inclusive. Grid Manager displays the **RFC 2317 Prefix** field. Enter a prefix in the text field. Prefixes can include alphanumeric characters. For information, see [Specifying an RFC 2317 Prefix](#) above.
    - **Name:** Enter the domain name of the reverse-mapping zone.
  - **DNS View:** This field displays only when there is more than one DNS view in the current network view. Select a DNS view from the drop-down list.
  - **Comment:** Optionally, enter additional information about the zone.
  - **Disable this zone:** Select this option to temporarily disable this zone. For information, see [Enabling and Disabling Zones](#). Note that disabling a zone may take a longer time to complete depending on the size of the data.
  - **Lock this zone:** Select this option to lock the zone so that you can make changes to it and prevent others from making conflicting changes. For information, see [Locking and Unlocking Zones](#).
4. Save the configuration, or click **Next** to continue to the next steps in the wizard as follows:
  - Define the name servers for the zone. For information on specifying primary and secondary servers, see [Assigning Zone Authority to Name Servers](#). For information on specifying authoritative name server groups, see [Using Authoritative Name Server Groups](#).

- If you have assigned a Microsoft server as the primary server for the zone and if the zone is AD-integrated, you can configure a list of domain controllers that are allowed to add NS records to the zone. For information see, [Configuring Domain Controller List](#).
- Define extensible attributes. For information, see [Using Extensible Attributes](#).

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

5. Click **Restart** if it appears at the top of the screen.

## Adding an IPv6 Reverse-Mapping Zone

To add an IPv6 reverse-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Auth Zone**.
2. In the *Add Authoritative Zone* wizard, click **Add an authoritative IPv6 reverse-mapping zone** and click **Next**.
3. Enter the following zone information:
  - Enter one of the following to identify the zone:
    - **IPv6 Network Prefix:** Enter the 128-bit IPv6 address for the address space for which you want to define the reverse-mapping zone. When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2006:0000:0000:0123:4567:89ab:0000:cdef can be shortened to 2006::123:4567:89ab:0:cdef. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The NIOS appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered. Choose the network prefix that defines the IPv6 network address space.
    - **Name:** Enter the domain name of the reverse-mapping zone.
  - **DNS View:** This field displays only when there is more than one DNS view in the current network view. Select a DNS view from the drop-down list.
  - **Comment:** Enter a descriptive comment about the zone.
  - **Disable:** Click this checkbox to temporarily disable this zone. For information, see [Enabling and Disabling Zones](#). Note that disabling a zone may take a longer time to complete depending on the size of the data.
  - **Lock:** Click this checkbox to lock the zone so that you can make changes to it and prevent others from making conflicting changes. For information, see [Locking and Unlocking Zones](#).
4. Save the configuration, or click **Next** to continue to the next steps in the wizard as follows:
  - Define the name servers for the zone. For information on specifying primary and secondary servers, see [Assigning Zone Authority to Name Servers](#). For information on specifying authoritative name server groups, see [Using Authoritative Name Server Groups](#).
  - Define extensible attributes. For information, see [Using Extensible Attributes](#).

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Adding an Authoritative Subzone

After creating a zone, you can add more zones at the same level, or add subordinate zones (*subzones*). The subzones can be authoritative, delegated, forward, or stub. For simplicity, the zones created in this example are authoritative (as are all zones by default). For information about configuring the other zone types, see [Configuring Delegated, Forward, and Stub Zones](#).

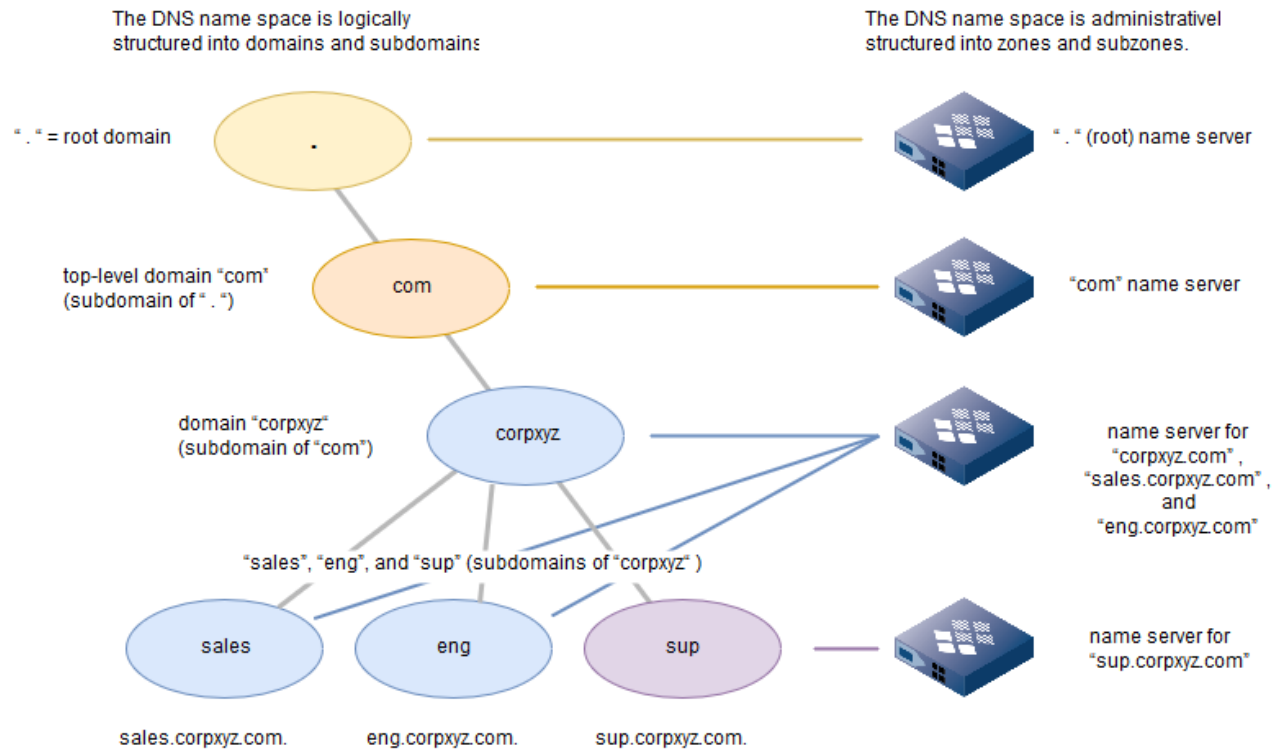
You create an authoritative zone when you assign authority for all the resource records of a particular domain to one or more name servers. You create a subzone when you assign authority for all the resource records of a subdomain to name servers. The name servers can be the same as, or different from, the name servers that serve resource records for the parent domain.


The distinction between domains and zones is that domains provide a logical structure to the DNS name space while zones provide an administrative structure. The difference between domains and subdomains and zones and subzones is

that the terms *subdomains* and *subzones* reference their relationship to a parent domain or zone. With the exception of the root domain and root zone, all domains are subdomains and all zones are subzones.

You can organize a domain based on logical divisions such as type (.com, .gov, .edu; or sales, eng, sup) or location (.uk, .jp, .us; or hq, east, west). The below figure shows one way to organize the external (public) name space and the internal (private) name space for a corporation with the domain name *corpxyz.com*. The external name space follows standard DNS conventions. Internally, you create an individual subdomain and corresponding subzone for each department.

*Domains and Subdomains, and Forward-Mapping Zones and Subzones*



 **Note**  
Throughout this documentation, the trailing dot indicating the root zone is not shown, although its presence is assumed.

The procedure for adding a subzone is the same as that used to add an authoritative zone. The only difference is that you specify the subzone name in the **Name** field. For information about adding authoritative zones, see [Configuring Authoritative Zones](#).

## Locking, Unlocking, Enabling, Disabling Zones

You can lock a zone when you create or edit it to prevent other administrators from making conflicting changes.

### Locking and Unlocking Zones

When you lock a zone, Grid Manager displays **LOCKED** beside the zone name when you view the records and subzones of the zone in the Zones panel. When other administrators try to make changes to a locked zone, the system displays a warning message that the zone is locked by *admin\_name*.

You can perform dynamic updates through mechanisms such as DDNS and nsupdate on a locked zone. The system can also add auto-generated records such as glue A records and NS records to a locked zone. Locks on a zone do not



impact its child zones.

Only a superuser or the administrator who locked the zone can unlock it. Locks do not expire; you must manually unlock a locked zone.

## Enabling and Disabling Zones

The NIOS appliance allows you to disable and enable a zone when you create or edit it. When you disable a zone, Grid Manager removes it from the DNS configuration file, but not from the database. This feature is especially helpful when you have to move or repair the server for a particular zone. You can easily disable a zone temporarily, and then enable it after the move or repair is completed.



### Notes

- When you temporarily disable a zone that has an associated NS group, the appliance removes all the automatically generated NS records, glue A or AAAA records, and PTR records from the zone. The appliance automatically generates the NS records, glue A or AAAA records, and PTR records when you re-enable the zone. Note that disabling a zone may take a longer time to complete depending on the size of the data.
- Do not enable authoritative zones if your Grid members have smaller disk spaces and if you want to perform DDNS or other updates on the authoritative zone. Infoblox recommends that you have a disk space of 250 GB if you want to use authoritative zones with Grid members.

## Configuring Authoritative Zone Properties

A zone inherits some of its properties from the Grid or from the member that serves it as a primary or secondary server. When you edit a zone, you can override properties set at the Grid or member level and modify the original zone settings, as well.

To configure authoritative zone properties:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox, and then click the Edit icon.
2. In the *Authoritative Zone* editor, you can do the following in each tab:
  - **General:** Modify the original zone settings, except the zone name.
  - **Name Servers:** Specify primary and secondary servers as described in [Assigning Zone Authority to Name Servers](#).
  - **Settings:** Set certain properties if the primary server is a Grid member. If the zone's primary server is an external server, then all these fields, except **Don't use forwarders to resolve queries in subzones**, are read-only with the information derived from the SOA record of the zone.
    - The **Serial Number** field displays the zone's current serial number. You can change the serial number in an SOA record only if the primary server of the zone is a Grid member. The serial number in the SOA record increments every time the record is modified. This serial number plays a key role when and how zone data is updated via zone transfers. The NIOS appliance allows you to change the serial number (in the SOA record) for the primary server so it is higher than the secondary server, thereby ensuring zone transfers come from the primary server (as they should).
    - Set **Refresh** intervals above 300 seconds, as Refresh intervals below 300 seconds may not work.
    - Override the Grid or member TTL settings as described in [About Time To Live Settings](#).
    - Override the email settings, as described in [Adding an Email Address to the SOA Record](#).
    - Change the primary name server that is specified in the SOA MNAME of a zone, as described in [Changing the SOA Name for a Zone](#).
    - **Don't use forwarders to resolve queries in subzones** : If the DNS members are configured to use forwarders to resolve queries that they cannot resolve locally, you can select this checkbox to disable the use of forwarders to resolve queries for data in the subzones.
  - **Queries:** Set restrictions for queries as described in [Controlling DNS Queries](#).



- **Zone Transfers:** Specify to which servers zone transfers are allowed as described in [Enabling Zone Transfers](#).
  - **Updates:** Set dynamic DNS update properties as described in [Configuring DNS Servers for DDNS](#).
  - **Active Directory:** Set parameters to allow zones to receive GSS-TSIG authenticated DDNS updates from DHCP clients and servers in an AD domain. For information, see [Supporting Active Directory](#).
  - **Extensible Attributes:** Define extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions:** Define administrative permissions. For information, see [About Administrative Permissions](#).
3. Click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, you can do the following in each tab:
    - **General:** Click the **Advanced** subtab and view the networks associated with the zone. This tab is visible only if the primary server is a Grid member, a Microsoft server, or unassigned. If a zone is associated with one or more networks, the IP addresses of its host, A and AAAA records must belong to the associated networks. You cannot change the network associations in this editor. Navigate to the *DHCP Network* editor of the network, to change the zone associations. For information, see [Associating Networks with Zones](#).  
You can enable fixed RRset ordering for the authoritative zone to save the order of the NAPTR records that are added to the zone using CSV import. Select the **Enable fixed RRset ordering** checkbox to enable fixed RRset ordering for the NAPTR records that are added to the authoritative zone through CSV import. There are a few best practices you should consider when enabling this feature. For more information, see [Enabling Fixed RRset Ordering for NAPTR Records](#).
    - **DNS Integrity Check:** Configure the appliance to monitor DNS data in the NS RRsets for authoritative zones. The appliance generates alerts when data discrepancies have been detected so you can mitigate possible DNS domain hijacking. For more information, see [Configuring DNS Integrity Check for Authoritative Zones](#).
    - **Host Naming:** Set restrictions for host names. For information, see [Specifying Hostname Policies](#).
    - **Shared Record Groups:** Add shared record groups to a zone. For information, see [Configuring Shared Record Groups](#).
    - **DNSSEC:** Configure DNSSEC properties. For information, see [Configuring DNSSEC](#).
    - **Record Scavenging:** Define the rules for DNS records scavenging in the zone, as described in [DNS Record Scavenging](#).
    - **Updates:** Click the **Advanced** subtab and specify the secure dynamic updates settings, as described in [Secure Dynamic Updates](#).
  4. Save the configuration and click **Restart** if it appears at the top of the screen.  
or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Enabling Fixed RRset Ordering for NAPTR Records

You can enable fixed RRset ordering for the authoritative zone to save the order of the NAPTR records that are imported to the zone using CSV import. When you enable fixed RRset ordering for the zone, the NAPTR records are saved in the same order as received in the CSV file and the appliance returns the NAPTR records in the fixed order in response to a query. If you enable fixed RRset ordering for an authoritative zone, the zone data is read-only and you cannot add, delete, or modify DNS records of the zone using GUI or API. You can add, delete, or modify DNS records of the zone using CSV import only. The fixed RRset ordering for the NAPTR records is retained only when the NAPTR records are imported to the zone through CSV import. If you change the position of the NAPTR records by disabling fixed RRset ordering, then the changes made for the position of the records will not be retained when you later enable fixed RRset ordering.

Note the following:

- When you import NAPTR records to an empty zone, you must use the **Add** CSV import option and to update the zone data you must use the **Replace** CSV import option to ensure that the RRset order of the NAPTR records is the same as specified in the CSV import file. For more information about CSV import feature, see [Guidelines for CSV Import](#).

- It is not recommended to configure DDNS updates for a zone, if you have enabled fixed RRset ordering for the zone.
- If you enable fixed RRset ordering for an authoritative zone with existing NAPTR records, the RRset order of the existing records will not change.
- The RRset order for NAPTR records imported to the zone through DNS zone transfers or DIW (Data Import Wizard) will not be preserved.
- You cannot enable fixed RRset ordering for an authoritative zone if the zone has subzones.
- You cannot enable fixed RRset ordering for an authoritative zone, if the zone has bulk hosts or if the zone is associated with a shared record group.
- This feature is applicable only for forward-mapping authoritative zones.
- This feature is not supported for Microsoft zones.

To enable fixed RRset ordering for NAPTR records imported to the zone using CSV import:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox, and then click the Edit icon.
2. In the *Authoritative Zone* editor, click **Toggle Advanced Mode** if the editor is in the basic mode.
3. Click the **Advanced** subtab of the **General** tab and then complete the following:
  - **Enable fixed RRset ordering:** Select this checkbox to enable fixed RRset ordering for the NAPTR records that are added to the authoritative zone through CSV import, as described in [Guidelines for CSV Import](#).
4. Save the configuration and restart DNS service for the change to take effect.

## Configuring DNS Integrity Check for Authoritative Zones

In certain DNS domain hijacking scenarios, hijackers alter the DNS data of a domain after gaining control of it. They consequently redirect users to a fraudulent site, instead of the legitimate site, on the Internet. To protect your authoritative DNS server against this type of DNS domain hijacking, you can configure the appliance to periodically monitor DNS data for top-level or parent authoritative zones. Based on your configuration, the appliance periodically checks DNS data in the NS RRsets for these zones and compares the data with that in the appliance database. It then reports data discrepancies through SNMP traps and logs related events in the syslog. You can also monitor the status of DNS data discrepancies, if any, through the *DNS Integrity Check* widget on the Task Dashboard. The severity in data discrepancies can help identify possible domain hijacking.

DNS integrity check is supported on all Infoblox appliances, including Advanced Appliances used primarily for Infoblox.

For information, see [About Infoblox Advanced DNS Protection](#). You can configure DNS integrity check for any selected authoritative zones, but you cannot configure it at the Grid, member, or DNS view level.

When you enable this feature, the appliance queries the NS RRsets and glue records for the top-level authoritative zones and compare the data with that in the appliance database. It does not query data for sub zones or delegated zones in the Grid.



### Note

DNS integrity check is not supported on authoritative zones configured to use primary DNS servers in stealth mode.

## Configuring DNS Integrity Check

To configure the appliance to check NS and glue records for a top-level or parent authoritative domain, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> top-level *authoritative zone* that you want to monitor, and then click the Edit icon from the Toolbar. Note that you can configure this feature only at the zone level. You can also configure zones that have the same name in different DNS views.  
Note that once you configure a zone for DNS integrity check, you will not be able to add a parent zone above this zone. You must disable DNS integrity check for this zone before you can add the parent zone.
2. In the *Authoritative Zone* editor, toggle to the **Advanced Mode**, select the **DNS Integrity Check** tab -> **Basic** tab and complete the following:

- **Enable:** Select this checkbox to enable the DNS integrity check feature.
  - **Member:** Click **Select Member** to select the Grid member you want to use for DNS integrity check. When you select a member, ensure that the member is configured to send and receive DNS queries and responses from Grid primaries (excluding stealth primaries) for the zone being monitored. Note that queries generated by DNS integrity check for the first reachable internal Grid primary are logged in relevant DNS reports. For information about reports, see [Infoblox Reporting and Analytics](#).
  - **Check Frequency:** Enter how often the appliance monitors DNS data for the authoritative zone. Select the time unit from the drop-down list. The appliance periodically queries DNS data for the top-level zone based on the time interval you configure here. The default value is one hour, and the minimum configurable value is 15 minutes.
  - **Enable Verbose Logging:** Select this to enable detailed logging of events related to DNS integrity check. When you select this option, the appliance logs additional information in the syslog when DNS data discrepancies are detected. It also logs a message when no data discrepancies are found during a DNS data check. When you clear this checkbox, the appliance logs standard information in the syslog and does not log an event when no data discrepancies are found during a DNS integrity check. This is disabled by default. For information about the syslog, see [Viewing the Syslog](#).
3. Save the configuration.

## Monitoring DNS Data Discrepancies for Authoritative Zones

When the appliance detects DNS data discrepancies between the authoritative and delegated zones, it reports the discrepancies through SNMP traps and email notifications, if configured. For more information, see [Setting SNMP and Email Notifications](#). The appliance classifies data discrepancies by severity, as follows:

- **Critical:** Data in the NS RRsets for the authoritative and delegate zones are completely out of synchronization.
- **Severe:** Some data in the NS RRset between the authoritative and delegate zones overlaps and some data is different.
- **Warning:** The NS RRset for the authoritative zone is a subset of the NS RRset for the delegate zone. It is possible that incorrect IP addresses have been entered at the registrar.
- **Informational:** The NS RRset for the delegate zone is a subset of the NS RRset for the authoritative zone. This could indicate a possible delay in domain registration.
- **Normal:** There are no DNS data discrepancies between the NS RRsets for the authoritative and delegated zones.
- **None:** No DNS discrepancies data has been collected or DNS integrity check has not been performed.

When different Grid primaries report different severity levels for the same data check, the appliance reports the most severe discrepancy level. When different Grid primaries report the same severity for the data check, the appliance reports only the first check.

You can use the following methods to monitor DNS data discrepancies for selected authoritative zones:

- Viewing syslog events, as described in [Viewing the Syslog](#).
- Monitoring DNS data discrepancy status through the *DNS Integrity Check* dashboard widget, as described in [DNS Integrity Check](#).
- Receiving SNMP traps and email notifications, as described in [Setting SNMP and Email Notifications](#).

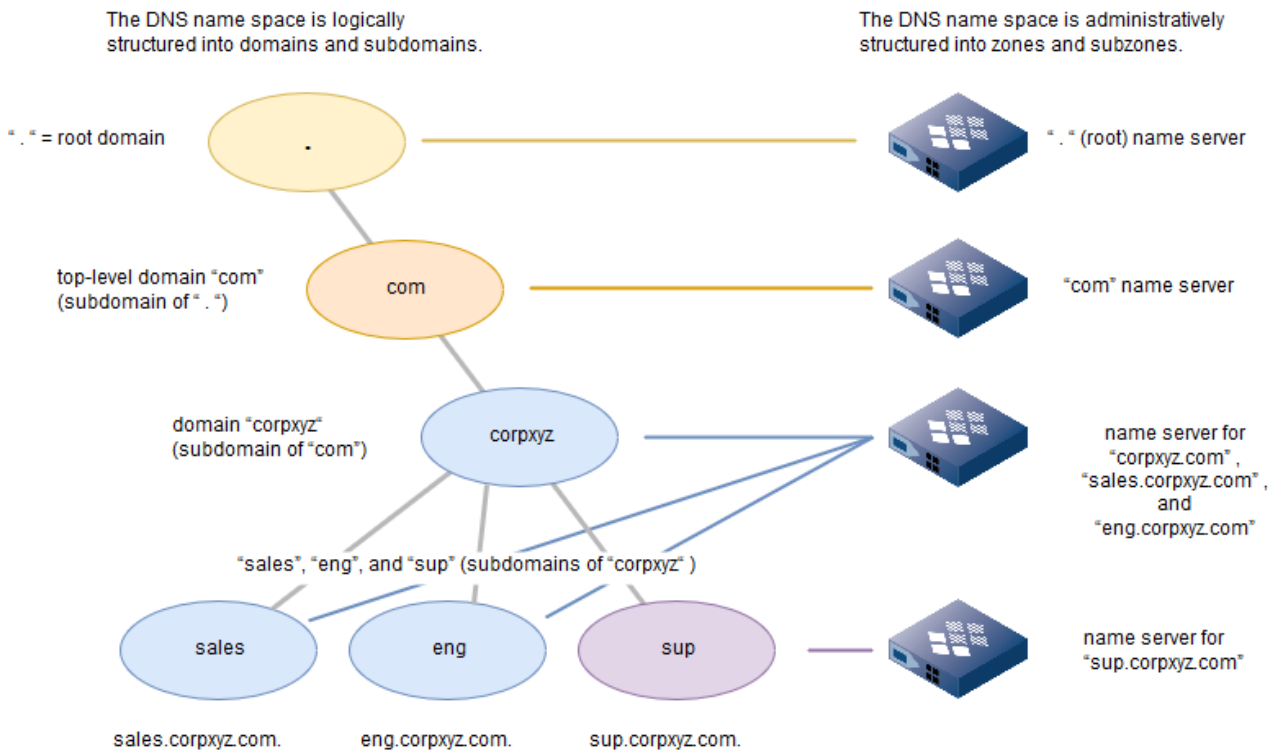
## Domains and Zones

After creating a zone, you can add more zones at the same level, or add subordinate zones (*subzones*). The subzones can be authoritative, delegated, forward, or stub.

The distinction between domains and zones is that domains provide a logical structure to the DNS name space while zones provide an administrative structure. The difference between domains and subdomains and zones and subzones is that the terms *subdomains* and *subzones* reference their relationship to a parent domain or zone. With the exception of the root domain and root zone, all domains are subdomains and all zones are subzones.

You can organize a domain based on logical divisions such as type (.com, .gov, .edu; or sales, eng, sup) or location (.uk, .jp, .us; or hq, east, west). The following figure shows one way to organize the external (public) name space and the internal (private) name space for a corporation with the domain name *corpxyz.com*. The external name space follows standard DNS conventions. Internally, you create an individual subdomain and corresponding subzone for each department.

*Domains and Subdomains, and Forward-Mapping Zones and Subzones*



On the Infoblox appliance, you can configure and manage DNS zones and subzones.

### IDN Support For DNS Zones

Grid Manager supports IDNs for DNS zones and resource records. For information about IDN, see [Managing Internationalized Domain Names](#). You can use either IDN or punycode (representation of IDN) to create DNS zones. Even if you use punycode to create a zone, the appliance automatically generates the corresponding IDN and displays the zone name in its native characters. Make sure that you use valid punycode to create a DNS zone. If you specify an invalid punycode, the appliance retains the punycode and does not convert it into IDN. Note that the appliance displays both the IDN and punycode for an IDN zone.

The following table summarizes how the appliance displays IDNs at the DNS zone level:

Input	NIOS Displays...	NIOS DNS Domain (Punycode in the GUI)	Conversion Guidelines
hello.com	hello.com	hello.com	No conversion
прывітанне.com	прывітанне.com	xn--80adk5aaihr3f9e.com	IDN to punycode
xn--80adk5aaihr3f9e.com	прывітанне.com	xn--80adk5aaihr3f9e.com	Punycode to IDN
\xyz format	\xyz format	\xyz format	No conversion

### Assigning Zone Authority to Name Servers

Forward-mapping zones answer name-to-address queries, and reverse-mapping zones answer address-to-name queries. When you create an authoritative forward-mapping zone or reverse-mapping zone, you assign zone authority to a name server and define it as the primary server for the zone. A primary server is designated as the primary source for

the zone and maintains a master copy of the zone data.

In a traditional DNS configuration, zone updates are performed based on a model that consists of a single primary server and one or multiple secondary servers, which receive read-only zone updates from the primary through database replications or zone transfers. Since the primary server contains editable zone data and is designated as the primary source for the zone, it can become a single point of failure when it becomes unavailable. To avoid a single point of failure for zone transfers, you can configure multiple primary servers for an authoritative zone.

When you define multiple primary servers for a zone, each primary has a copy of the zone's authoritative data that can be updated independently. When you modify zone data, the appliance replicates updated data among all primary servers. When there are conflicts between zone updates, the appliance generally selects the latest updates based on the timestamps the updates were made by the clients to the primary servers. Therefore, accurate time synchronization among all servers in the DNS configuration is very important. For more information about other best practices for configuring multiple primaries for an authoritative zone, see [Best Practices for Defining Multiple Primaries for Authoritative Zones](#) below.

You can also create one or more secondary name servers for a zone. A secondary server for a zone receives read-only zone data from the primary server. If a zone is part of an internal DNS structure for a private network, the inclusion of a secondary DNS server is optional, though highly recommended. If a zone is part of an external DNS structure for a public network such as the Internet, then a secondary server in a different subnet from the primary server is required. This requirement provides an additional safeguard against localized network failures causing both primary and secondary name servers for a zone to become inaccessible.



#### Note

You can enter, modify, and remove zone data on the primary servers, which can then send new and modified data in a read-only format to the secondary servers. Both primary and secondary name servers are authoritative for the zone data they serve. The distinction between them is how they get their zone data.

In Grid Manager, you can specify the primary and secondary servers for a zone or you can specify a name server group. A name server group is a collection of one or more primary servers and one or more secondary servers. For information on name server groups, see [Using Name Server Groups](#).

## Best Practices for Defining Multiple Primaries for Authoritative Zones

Before you configure multiple primary servers for a zone, consider the following guidelines to ensure data integrity:

- This feature is designed to increase availability of the DNS service by allowing multiple primaries for a zone. It will not increase overall throughput of DNS update traffic, as ultimately all updates must be replicated to (and processed by) all of the primaries.
- When determining which appliances should act as primaries for the zone, consider that an additional SOA record will be required in the database for each primary. This will add to the overall record count for the zone, and each SOA will need to be updated for any change to the zone, which can impact performance.
- Enable NTP for all members (at the member level) and ensure that their times are properly synchronized with their local time servers. Ensure that you select the "Exclude the Grid Master as an NTP server" option. The appliance selects the latest zone updates based on the timestamps the updates were made by clients to the primary servers. This is especially important when there are conflicts between two or more zone updates. For information about NTP, see [Using NTP for Time Settings](#).
- When specifying the primary server for secondaries, you can choose to have the appliance automatically select it for you based on latency determination or you can manually specify it. When manually selecting a primary for zone updates, consider using one that is close in proximity to the secondary servers, which can result in better service performance. For information about setting preference for the primary server, see [Adding Grid Secondaries](#) below.
- You can configure a default primary for DDNS updates to a zone with multiple primary servers. To enhance service performance, select a default primary that is close in proximity to the DHCP server that provides DDNS updates. This is especially useful if you have DHCP members that are located in different locations. You can configure a different default primary for each DHCP member based on their locations. For more information, see [Defining the Default Primary for DDNS Updates to Zones with Multiple Primaries](#).
- DNSSEC is not supported for zones with multiple primary servers. These zones must be unsigned. For information about DNSSEC, see [Configuring DNSSEC](#).

- When determining which appliances should act as primaries for the zone, consider that an additional SOA record will be required in the database for each primary. This will add to the overall record count for the zone, and each SOA will need to be updated for any change to the zone, which can impact performance.

## Specifying a Primary Server

When you create a zone, the primary server can be a Grid member, an external DNS server that you specify, or a Microsoft DNS server that is managed by a Grid member. For information about managing Microsoft Windows DNS servers, see [About Managing Microsoft Windows Servers](#).

Although a zone typically has only one primary server, you can specify multiple primary servers for an authoritative zone. You can configure multiple Grid primaries or multiple external primaries (including Microsoft AD-integrated servers) for a zone, but you cannot configure both at the same time for the same zone. In addition, you can configure one Microsoft server, but not multiple Microsoft servers (except for Microsoft AD-integrated servers), as the primary server for a zone. Note that each primary server that you configure for a zone has its own MNAME for the SOA record and serial number. For information about how to view and modify certain values in the SOA record, see [Viewing and Modifying SOA Records](#).

A hidden primary provides data to its secondary servers, which in turn respond to DNS queries using this data. One of several advantages of this approach is that you can take the primary server offline for administrative or maintenance reasons without causing a disruption to DNS service (within the expiration interval set for the validity of its zone data—the default is 30 days).

When you add an authoritative forward-mapping zone and assign responsibility for the zone to a primary name server whose host name belongs to the name space of the zone, the NIOS appliance automatically generates an NS (name server) record and an A (address) record for the name server. This type of A record is called a glue record because it "glues" the NS record to the IP address (in the A record) of the name server.

In Grid Manager, you can specify the primary server for a zone when you create it using the *Add Authoritative Zone* wizard or when you edit an existing zone using the *Authoritative Zone* editor. For information on how to add a new zone through the wizard, see [Configuring Authoritative Zones](#). The following procedure describes how to access the editor of a zone. To specify a primary server for an existing zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox, and then click the Edit icon.
2. In the *Authoritative Zone* editor, click **Name Servers**.
3. Select **Use this set of name servers**.
4. Click the Add icon and select one of the following options for a primary server:
  - **Grid Primary:** Choose this option to select a Grid member as the primary server for the zone. See [Specifying Grid Primary Servers](#) below.
  - **Microsoft Primary:** Choose this option to select a Microsoft DNS server as the primary server for the zone. See [Specifying Microsoft Primary Servers](#) below.
  - **External Primary:** Choose this option if the appliance is in a Grid and you want to specify a primary server outside the Grid ("external" to the Grid). See [Specifying External Primary Servers](#) below.
5. Save the configuration and click **Restart** if it appears at the top of the screen. or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Specifying Grid Primary Servers

In the *Add Grid Primary* panel, do the following, and then click **Add** to add the Grid member to the list of name servers for the zone as primary:

- If no member is displayed, click **Select** to specify a Grid member. When there are multiple members, Grid Manager displays the *Member Selector* dialog box from which you can select a primary name server. To select multiple primary servers, click **Select** and then **Add** again.
- **Stealth:** Click this to hide the NS record for the primary name server from DNS queries. The NIOS appliance does not create an NS record for the primary name server in the zone data. Clear the checkbox to display the NS record for the primary name server in responses to queries.

## Changing the SOA Name for a Zone

If the primary name server of a zone is a Grid member, the NIOS appliance allows you to change the SOA (start of authority) name that is automatically created when you initially configure the zone. For example, you might want to hide the primary server for a zone. If your appliance is named dns1.zone.tld, and for security reasons, you may want to show



a secondary server called dns2.zone.tld as the primary server. To do so, you would go to dns1.zone.tld zone (being the true primary) and change the SOA to dns2.zone.tld to hide the true identity of the real primary server.

To change the SOA name for a zone:

1. From the **Data Management** tab, select the **DNS** tab > **Zones** tab > *dns\_view* -> *zone* checkbox -> Edit icon.
2. In the *Authoritative Zone* editor, click **Settings**.
3. Click **Override** beside the **Primary name server** field and enter the new SOA name. This field supports IDN.
4. Save the configuration and click **Restart** if it appears at the top of the screen.  
or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

### Specifying Microsoft Primary Servers

You can assign a Microsoft server as the primary server of a zone when it is managed by a Grid member in read/write mode. For information, see [About Managing Microsoft Windows Servers](#). When a Microsoft server is the primary server of a zone, the zone supports only standard DNS resource records. It does not support the Infoblox record types host records, bulk host records, and shared record groups. You cannot add any of these records to the zone nor assign a DNS zone with these records to a Microsoft server as the primary server.

In the Add Grid Primary panel, do the following to assign a Microsoft primary server:

1. Complete the following:
  - Select **Use this set of name servers**.
  - Click the *Add* icon and select **Microsoft Primary**.
2. In the *Add Microsoft Primary* panel, do the following, and then click **Add** to add the Microsoft primary server to the list of name servers for the zone:
  - If no server is displayed, click **Select Server** to specify a Microsoft server. When there are multiple servers, Grid Manager displays the *Server Selector* dialog box from which you can select a Microsoft server. Grid Manager lists Microsoft servers that are managed in read/write mode. It does not include Microsoft servers managed in read-only mode.
  - **Information to create NS record:** Grid Manager automatically creates the NS record. After you select a server, Grid Manager populates the **Name** and **IP Address** fields. Grid Manager uses this information when it creates the NS record, unless you select **Stealth**. You can specify a different FQDN or IP address for the NS record; for example, for a multihomed server.
  - Store the zone in Active Directory (AD Integrated Zone): This is enabled and selected by default only if the Microsoft server is a domain controller. Note that you can enable Active Directory integration only after the Microsoft server has been synchronized at least once because its AD ability is not known before the synchronization. This is disabled when the Microsoft server is not a domain controller.
  - **Stealth:** Select this option to hide the NS record for the primary name server from DNS queries. Grid Manager does not create an NS record for the primary name server in the zone data. Clear this option to display the NS record for the primary name server in responses to queries. Note that this option is not available for AD-integrated zones.

### Specifying External Primary Servers

In the Add External Primary panel, do the following, and then click **Add** to add the external primary server to the list of name servers for the zone:

- **Name:** Type a resolvable domain name for the external primary server.
- **Address:** Type the IP address of the external primary server.
- **Multi-master:** This appears only when there is more than one external primary assigned to the zone. Select this checkbox for external primary servers when the zone is in another Grid and has multiple Grid primaries. When you select this option, it is selected for all external primaries assigned to the zone. This zone is identified as an external zone with multiple primary servers.
- **Use TSIG:** To authenticate zone transfers between the local appliance and the external primary server using a TSIG (transaction signature), select this checkbox. Infoblox TSIGs use HMAC-MD5 hashes. These are keyed one-way hashes for message authentication codes using the Message Digest 5 algorithm. For details, see *RFC 1321, The MD5 Message-Digest Algorithm*, and *RFC 2104, HMAC: Keyed-Hashing for Message Authentication*.

- **Key name:** Type or paste the name of the TSIG key you want to use. This must be the same name as that of the TSIG key on the external primary server.
- **Key Data:** Type or paste a previously generated key. This key must also be present on the external primary server. You can generate a TSIG key, or obtain the TSIG key name and key from the external name server, by accessing the server yourself or by requesting the server administrator to deliver them to you through some out-of-band mechanism. Then type or copy-and-paste the name and key into the appropriate fields.
- **Use 2.x TSIG:** If you want to use TSIG authentication and the external primary name server is a NIOS appliance running DNS One 2.x code, select this checkbox. The local appliance generates the required TSIG key for authenticating DNS messages to and from appliances running DNS One 2.x code.



#### Note

On the appliance you configure as a secondary server for a zone, you must associate a TSIG key for each primary server to which the secondary server requests zone transfers. On the appliance you configure as a primary server for a zone, you can set a TSIG key at the Grid, member, or zone level. Because the secondary server requests zone transfers, it must send a specific key in its requests to the primary server. Because the primary server responds to the requests, it can have a set of TSIG keys from which it can draw when responding. As long as the primary server can find the same TSIG key that the secondary sends it, it can verify the authenticity of the requests it receives and authenticate the responses it sends. Use NTP to synchronize the time on both name servers that use TSIG-authenticated zone transfers.

## Specifying Secondary Servers

A secondary name server is as authoritative for a zone as a primary server. Like a primary server, a secondary server answers queries from resolvers and other name servers. The main difference between a secondary and primary server is that a secondary server receives all its data from a primary server, or possibly from another secondary server that relays zone data it receives. The zone data passes from a primary to a secondary server (and possibly from that secondary server on to another secondary server). This process is called a zone transfer.

The advantage of using primary and secondary name servers is that you enter and maintain zone data in one place—on the primary server. The data is then distributed to the one or more secondary servers.

Secondary servers can be Grid members, external DNS servers or Microsoft DNS servers that are managed by Grid members. In Grid Manager, you can specify the secondary server for a zone when you create it using the *Add Authoritative Zone* wizard and when you edit an existing zone using the *Authoritative Zone* editor. For information on how to add a new zone through the wizard, see [Configuring Authoritative Zones](#). The following procedure describes how to access the editor of a zone.

To specify a secondary server for an existing zone:

1. From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *zone* checkbox, and then click the Edit icon.
2. In the *Authoritative Zone* editor, click **Name Servers**.
3. Select **Use this set of name servers**.
4. Click the Add icon and select one of the following options:
  - **Grid Secondary:** Selects the local appliance as the secondary server (or if the appliance is deployed in a Grid and you want to make a different member the secondary server). See [Adding Grid Secondaries](#) below.
  - **Microsoft Secondary:** Select this option if you want to specify a managed Microsoft DNS server as a secondary server. See [Specifying Microsoft Secondary Servers](#) below.
  - **External Secondary:** Select this option if the appliance is in a Grid and you want to specify a secondary server outside the Grid ("external" to the Grid), or if the appliance is deployed independently from a Grid. See [Specifying External Secondaries](#) below.
5. Save the configuration and click **Restart** if it appears at the top of the screen. or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

### Adding Grid Secondaries

When adding Grid secondaries to a zone that has multiple primary servers, the appliance selects a primary server as the active server based on the method that you have selected. If you select **Automatic**, the primary is selected based on latency determination, which occurs separately on each primary. When available, the primary server that has the lowest latency is preferred. When you select **Manual**, latency determination is ignored, and the first available primary server in



the list is selected as the active server. Thus if the first primary on the list is not available, the next available primary is used. Depending on which primary server is selected, the Grid secondary returns the FQDN of the primary in the MNAME field of the zone SOA record. It also includes the version of the zone content that it serves.

In the Add Grid Secondary panel, enter the following, and then click **Add** to add the Grid secondary server to the list of name servers for the zone:

- If no member is displayed, click **Select** to specify a Grid member. When there are multiple members, Grid Manager displays the *Member Selector* dialog box from which you can select a secondary name server.
- **Stealth:** This setting applies only if the primary server is a Grid member or Microsoft server. Select this to hide the NS record for the secondary name server from DNS queries. The NIOS appliance does not create an NS record for this name server in the zone data. Select the checkbox again to display the NS record for the secondary name server in responses to queries. A secondary server in stealth mode is also known as a "hidden secondary". For example, you can configure a hidden secondary when a secondary server is at a branch office with a slow connection to the rest of corporate network. Configure local hosts at the branch office to send DNS queries to the secondary server, but keep it hidden from other name servers on the rest of the network so that they do not send it queries. Instead, they use a server located in a different part of the network that has faster connection speeds.
- **Lead Secondary:** This option becomes available only after you specify the primary name server as external. When a primary server is external to a Grid whose members are secondary servers, you can select this checkbox to designate one member as a lead secondary. The primary server sends zone transfers to the lead secondary, which distributes the zone data to the other secondary servers in the Grid using zone transfers (not the Grid data replication mechanism). After you designate a Grid member as a lead secondary for a zone, you do not have to configure members to use the lead secondary server. All other Grid members acting as secondary servers for the zone automatically use the lead secondary to get zone data. Using a lead secondary simplifies the addition, modification, and removal of other secondary servers in the Grid. As long as the lead secondary remains unchanged, you need not update intervening firewall policies or the external primary server whenever you make changes to non-lead secondary Grid members. This approach also reduces the amount of traffic between primary and secondary servers.
- **Update Zones Using:** This option becomes available only after you specify a Grid member as the primary server.
  - **Grid Replication (recommended):** Select this checkbox to use Grid replication to move zone data from the primary to secondary servers.
  - **DNS Zone Transfers:** Select this checkbox to use the DNS zone transfer process to move zone data from the primary to secondary servers.
- **Primary Preference:** This appears only for Grid secondaries that are assigned to zones with multiple primaries. Select one of the following to set preference for selecting the primary server:
  - **Automatic:** Choose this to allow the appliance to select the optimum primary server for zone updates based on latency determination, which prefers the primary server that has the lowest latency. This is selected by default.
  - **Manual:** Choose this to manually select primaries for zone updates. In the Available Primaries and Selected Primaries tables, click a name server and use the arrows to select and deselect primaries between the tables. Then use the up and down arrows next to the Selected Primaries table to put the selected primaries in a preferred order. The secondary servers will get zone updates from the selected primary server based on the given order. To optimize service performance, select primary servers that are close in proximity to the secondary servers.

### Specifying Microsoft Secondary Servers

You can assign a Microsoft server as the primary server of a zone when it is managed by a Grid member in read/write mode. For information, see [About Managing Microsoft Windows Servers](#).

Since Microsoft servers cannot replicate data from the Grid, when a DNS zone is defined as a secondary on a Microsoft server, the Microsoft server obtains the content of the zone only through DNS zone transfers.

- In the Add Microsoft Secondary panel, do the following:
  - If no server is displayed, click **Select Server** to specify a Microsoft server. When there are multiple servers, Grid Manager displays the *Server Selector* dialog box from which you can select a Microsoft server. Grid Manager lists Microsoft servers that are managed in read/write mode. It does not include Microsoft servers managed in read-only mode.
  - **Information to create NS record:** Grid Manager automatically creates the NS record. After you select a server, Grid Manager populates the **Name** and **IP Address** fields. Grid Manager uses this information when it creates the NS record, unless you select Stealth.

- **Stealth:** This setting applies only if the primary server is a Grid member or a Microsoft server. Select this option to hide the NS record for the secondary name server from DNS queries. Grid Manager does not create an NS record for this name server in the zone data. Clear this option to display the NS record for the secondary name server in responses to queries.

## Specifying External Secondaries

In the Add External Secondary panel, enter the following, and then click **Add** to add the external secondary server to the list of name servers for the zone:

- **Name:** Enter a resolvable domain name for the external secondary server.
- **Address:** Enter the IP address of the external secondary server.
- **Stealth:** This setting applies only if the primary server is a Grid member or a Microsoft server. Click this checkbox to hide the NS record for the secondary name server from DNS queries. The NIOS appliance does not create an NS record for the secondary name server in the zone data. Select the checkbox again to display the NS record for the secondary name server in response to queries.  
Note that to avoid an impact on your database performance, Infoblox recommends that you do not configure a large number of external secondary servers in stealth mode. To ensure that these secondary servers receive notifications about zone updates, you can allow zone transfers for these IP addresses and then enable the appliance to add them to the also-notify statement. For information about how to configure this feature, see [Configuring Zone Transfers](#).
- **Use TSIG:** To authenticate zone transfers between the local appliance and the external secondary server using a TSIG (transaction signature), select this checkbox. Infoblox TSIGs use HMAC-MD5 hashes. These are keyed one-way hashes for message authentication codes using the Message Digest 5 algorithm. For details, see RFC 1321, *The MD5 Message-Digest Algorithm*, and RFC 2104, *HMAC: Keyed-Hashing for Message Authentication*.
- **Key name:** Type or paste the name of the TSIG key you want to use. This must be the same name as that of the TSIG key for this zone on the external secondary server.
- **Key:** Type or paste a previously generated key. On the external secondary server, this key must also be present and associated with this zone. You can generate a TSIG key, or you can obtain the TSIG key name and key from the external name server, either by accessing the appliance yourself or by requesting the appliance administrator to deliver them to you through some out-of-band mechanism. Then, type or copy-and-paste the name and key into the appropriate fields.
- **Use 2.x TSIG:** Select this checkbox to use TSIG authentication and the external secondary name server is a NIOS appliance running DNS One 2.x code. The local appliance generates the required TSIG key for authenticating DNS messages to and from appliances running DNS One 2.x code.



### Note

On the appliance you configure as a secondary server for a zone, you must associate a TSIG key for each primary server to which the secondary server requests zone transfers. On the appliance you configure as a primary server for a zone, you can set a TSIG key at the Grid, member, or zone level. Because the secondary server requests zone transfers, it must send a specific key in its requests to the primary server. Because the primary server responds to the requests, it can have a set of TSIG keys from which it can draw when responding. As long as the primary server can find the same TSIG key that the secondary sends it, it can verify the authenticity of the requests it receives and authenticate the responses it sends. Use NTP to synchronize the time on both name servers that use TSIG-authenticated zone transfers.

## Configuring Domain Controller List

When you configure an AD-integrated authoritative zone, by default all the domain controllers that belong to the domain automatically add NS records to the AD-integrated zone, which might be undesirable in some deployment scenarios. For example, if a domain controller is deployed in a company's branch office, it is unlikely that the domain controller should be registered as the name server for the company's top-level zones. By configuring a domain controller list for the AD-integrated zone, you can control which domain controllers are allowed to add NS records to the zone. If you configure a domain controller list for an AD-integrated zone, only those in the list can add NS records to the zone. You can configure the domain controller list for AD-integrated zones either on the NIOS appliance or on the Microsoft server.

## Configuring Domain Controller List on NIOS

You can configure a domain controller list for an AD-integrated authoritative zone to allow NS record creation for specific domain controllers. You can add, modify, and delete the entries in the domain controller list if the Microsoft server assigned to the zone is managed in read/write mode. If you have not configured the domain controller list while configuring an AD-integrated zone, you can perform the configuration later while editing the zone.

Note the following about the domain controller list:

- If the domain controller list is empty for an AD-integrated zone, all domain controllers that belong to the domain can add NS records to the AD-integrated zone.
- If you remove a domain controller from the domain controller list of an AD-integrated zone, the NS record that belongs to the domain controller is not automatically deleted. You must manually delete the NS record from the AD-integrated zone.

To configure a domain controller list for an AD-integrated zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> click the Add icon and select **Authoritative Zone** from the drop-down menu.
2. Complete the details as mentioned in [Configuring Authoritative Zones](#) and then complete the following to specify a list of domain controllers that are allowed to add NS records to the AD-integrated zone:
  - **No**: Select **No** if you do not want to configure a domain controller list for the AD-integrated zone. If you select **No**, all domain controllers that belong to the domain can add NS records to the AD-integrated zone.
  - **Yes**: Select **Yes** to configure a domain controller list for the AD-integrated zone. If you select **Yes**, click the Add icon of the **Domain Controller** table and select one of the following from the drop-down menu:
    - **Add**: Select this to manually add the IP addresses of the domain controllers to the list. Grid Manager adds a new row to the table. Specify the following:
      - **Address**: Enter the IP address of the domain controller.
      - **Comment**: Enter information about the domain controller.To remove an entry from the list, select the entry and click the Delete icon.
    - **Auto-populate**: Select this to automatically populate the **Domain Controller** table with the list of domain controllers. In the *Add Prepopulated Domain Controllers* panel, select one of the following options:
      - **Zone**: Select this to copy the list of domain controllers from an existing AD-integrated zone in the NIOS database. Click **Select** to select the AD-integrated zone. Click **Clear** to remove the selected zone.
      - **Servers in Domain**: Select this to add the IP addresses of all the Microsoft servers available in the NIOS database which belong to the same AD domain as the primary Microsoft server assigned to the zone.
    - Click **Add** to add the list of domain controllers to the table. Grid Manager automatically populates the **Domain Controller** table with the list of domain controllers in ascending order by IP address. Note that the **Auto-populate** option to add the domain controller list is only available while configuring the AD-integrated zone. It is not available when you edit the domain controller list in the *Authoritative Zone* editor.
3. Save the configuration.

## Configuring Domain Controller List on the Microsoft Server

You can configure a list of domain controllers that are allowed to add NS records to an AD-integrated zone on the Microsoft server using the `dnscmd` command line utility as follows:

```
dnscmd DNS Server /Config Zone Name /AllowNSRecordsAutoCreation Ip Addresses
```

For example, if you are configuring a domain controller list for an AD-integrated zone `foo.net` on the DNS server, `192.69.0.1`, use the following command:

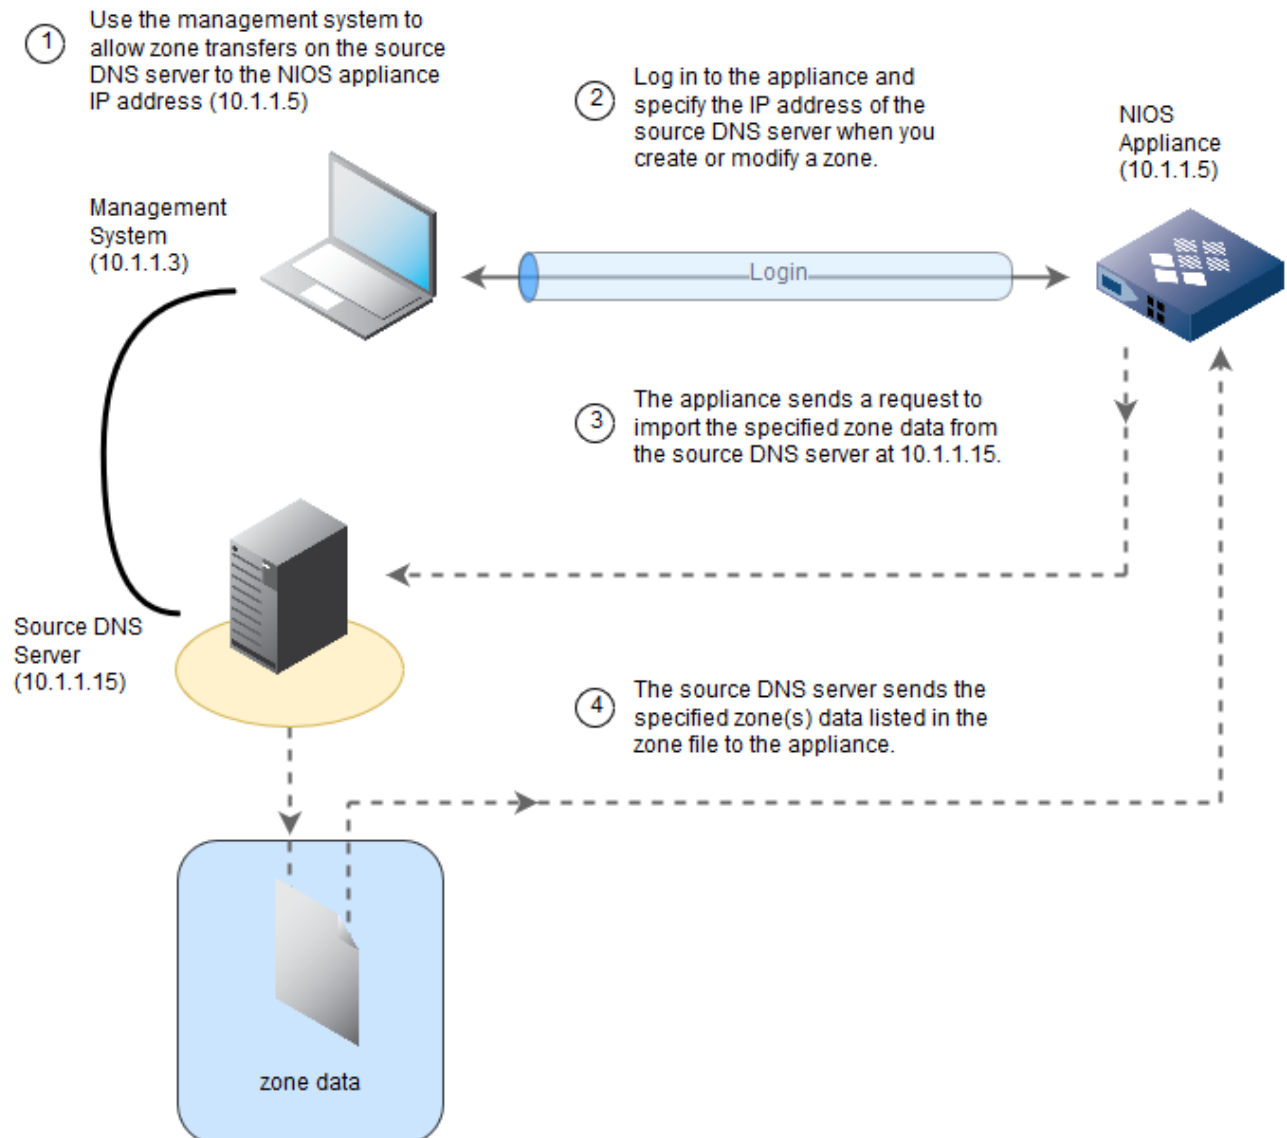
```
dnscmd 192.69.0.1 /config foo.net /AllowNSRecordsAutoCreation 192.69.0.6  
192.69.0.9
```

For more information about configuring a domain controller list for an AD-integrated zone on the Microsoft server see <https://technet.microsoft.com/en-us/library/cc755848%28v=ws.10%29.aspx>.

## Importing Zone Data

Importing zone information alleviates having to manually enter data through the Infoblox GUI. You can import data from existing name servers, as well as from NIOS appliances running version 3.1r4 or later. You can import existing zone data when you create a new zone and when you edit an existing zone. You can import one zone (and its subzones) at a time. For the remainder of this section, the name server that stores the existing zone data (which is imported) is referred to as the *source* name server (regardless of whether it is a third-party server or another NIOS appliance). The appliance that receives the zone data is referred to as the *destination* appliance. The following illustration shows the import zone data process.

### Importing Zone Data Process



The appliance imports zone data through a zone transfer. Therefore, the source name server must be authoritative for the zone data being imported. You must also configure the source name server to allow zone transfers to the destination appliance. On the source name server, you might need to modify the *allow-transfer* substatement to include the IP address of the destination appliance prior to importing the data. If you are importing zone data to an HA pair, use the VIP (virtual IP) address shared by the HA pair. For a single independent appliance, use the LAN IP address. If you are importing zone data to a Grid, always use the IP address of the Grid Master.

If the source name server is an Infoblox appliance, you can configure it to allow zone transfers as described in [Enabling](#)

[Zone Transfers](#). Note that a NIOS appliance, acting as the primary name server for a zone, by default allows zone transfers to its secondary name servers. If the zone import fails, the zone to which the data is imported will be disabled and the system does not create records and delegated subzones.



#### Note

The appliance does not encode punycode when you import zone data containing punycode. For example, a zone data containing IDNs in punycode is stored in punycode for the data being imported. The data is managed in punycode only.

### About Importing Data into a New Zone

When the appliance imports data to a newly created zone, it imports the existing A, CNAME, DNAME, SRV, TXT, MX, PTR, host, and bulk host records, but creates NS (and A records matching that NS record) and SOA records appropriate for the destination server. The NS and SOA records are auto-created when a destination appliance is specified as the primary or secondary name server for the new zone. If the imported zone has extra NS records, they are rewritten to specify the source server as an external secondary. Delegation is also added for any subzones. The subzone records are not imported.

### About Importing Data into an Existing Zone

When you import zone data into an existing zone, the zone retains the NS and SOA records automatically created when the zone was originally created and replaces all other records—A, PTR, MX, TXT, SRV, CNAME, DNAME, host, and bulk host. The local appliance also retains subzones and records in the subzones that exist locally. If there are no duplicates, the destination appliance records are retained. If the imported zone has extra NS records, those records change to designate the source server as an external secondary.

### Importing Data into Zones

In Grid Manager, you can import zone data when you create the zone using the *Add Authoritative Zone* wizard and when you edit an existing zone. For information on how to add a new zone through the wizard, see [Configuring Authoritative Zones](#). The last step of the wizard provides the option to import zone data. The following procedure describes how to import data into an existing zone.

**Note:** Only superusers can import zone data that contains A, AAAA, shared A, or shared AAAA records with a blank name. Limited-access users must have read/write permission to **Adding a blank A/AAAA record** in order to import zone data that contains A, AAAA, shared A, or shared AAAA records with a blank name, otherwise the import zone data operation might fail. You can assign global permission for specific admin groups and roles to allow to import A, AAAA, shared A, or shared AAAA records with a blank name. For more information, see [Administrative Permissions for Adding Blank A or AAAA Records](#).

To import data into an existing zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox, and then click **Import Zone** in the Toolbar.
2. In the *Import Zone* dialog box, specify the following:
  - The IP address of the name server from which you want to import data.
  - Optionally, click the **Automatically create Infoblox host records from A records** checkbox.
3. Click **Import**.

When the local server successfully imports the zone data, a *Confirmation* message appears. If the local server cannot import the zone data, an *Error* message appears, recommending that you verify the correctness of the IP address of the remote server and zone information.



#### Note

If NIOS resolves the IP address of the imported zone data, an external secondary member is added to the list of name servers with the exact IP address. If NIOS cannot resolve the IP address of the imported zone data, it adds an external secondary member with the IP address 255.255.255.255 to the list of name servers.

## Enabling Zone Transfers

A zone transfer is the process of sending zone data across a network from one name server to another. When the primary server detects a change to its zone data, it notifies the secondary servers. The secondary servers reply by checking to see if the serial number they have for the zone is as large as the serial number for the zone on the primary server. If not, the secondary servers request a zone transfer.

In addition to receiving zone change notifications, a secondary server periodically polls the primary server to see if their zone data is in sync. In response, the primary server can send a DNS message containing just the changed zone data, or the entire data set. The first type of transfer is known as an incremental zone transfer, or IXFR. The second type of transfer is known as a full zone transfer, or AXFR.

A NIOS appliance, acting as the primary name server for a zone, allows zone transfers to secondary name servers by default. This includes all servers listed in the NS records for that zone. (Secondary name servers in a Grid, however, receive updated zone data via database replication by default, as explained later in this section.) You can also specify zone transfers to other name servers, such as when migrating zone data to a new server or to a management system. You can specify one or more destinations to which the local appliance sends zone transfers. You can also specify the security and format of the transfers.

Note that secondary name servers periodically query the primary name server to find out if zone data has been changed. Each query takes a certain amount of time and bandwidth on the network. By default, secondary name servers limit the rate (serial-query-rate) at which these queries are being sent. Thus when the secondary name servers are serving a large number of zones, it may take a long time to detect changes to their zone data. You can configure this value to optimize the query rate on the network. In addition, when you have set up a few secondary name servers for a large number of zones, a delay in zone transfers may occur due to the default zone transfer configuration that limits concurrent zone transfers to 10 per secondary server. You can configure the maximum value of concurrent zone transfers to optimize the zone transfer operation. For information about how to optimize zone transfers, see [Configuring Concurrent Zone Transfers](#) below.

By default, Grid members automatically receive updated zone data via database replication (through an encrypted VPN tunnel). You can change the default behavior to allow Grid members to use zone transfers instead of Grid replication. Keep in mind that a database replication updates zone data for both the active and passive nodes of an HA member. Therefore, if there is a failover, the new active node (the previous passive node) immediately begins serving zone data with fresh information. In the case of a zone transfer, the passive node does not receive zone data until after a failover, when it becomes an HA master. At that time, it performs a zone transfer. If there is a lot of zone data, the transfer can take up to several minutes, thereby causing a break in the availability of the new HA master.

If you have HA members as secondary servers, zone transfers can result in service interruption when there is a failover. Furthermore, if the primary server is down when the HA member fails over, the new active node cannot receive zone data until the primary server comes back online.

You can use TSIG (transaction signature) keys to authenticate zone transfer requests and replies. The same key name and key value must be on the primary and secondary name servers for TSIG-authenticated zone transfers to occur. When using TSIG, it is important that both appliances involved with the authentication procedure use NTP (Network Time Protocol) for their time settings (see [Using NTP for Time Settings](#)).

You can control zone transfers at the Grid, member, and zone levels. This enables you to specify a different set of name servers for a Grid, member, and zone, if necessary. You can also control which external secondary servers should receive notifications about zone updates by adding their IP addresses to the also-notify statement for each authoritative zone that is served by a Grid member. Infoblox recommends that you use this feature to notify hidden external secondary servers about zone updates, instead of putting them in stealth mode, especially when you plan to configure a large number of them. For information about how to add IP addresses to the also-notify statement, see [Configuring Zone Transfers](#) below.



## Configuring Zone Transfers

To configure zone transfers, you identify the servers to which zone data is transferred and optionally, servers to which data must not be transferred. For example, you can allow transfers to a network, but not to a specific server in the network.

You can specify a different set of servers for specific Grid members and zones. For example, if certain Grid members are primary servers for a zone, then you can specify the secondary servers to which that member is allowed to transfer zones.

You can also enable the appliance to add all IPv4 and IPv6 addresses for which you allow zone transfers to the also-notify statement for each authoritative zone that is served by a Grid member. The also-notify statement defines a list of addresses that receive notifications about zone updates, in addition to the IP addresses listed in the NS records for the zone. Use this feature to configure a large number of external secondary servers, instead of putting them in stealth mode.

To configure zone transfer properties:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab, click the **Members** tab -> *member* checkbox -> Edit icon.  
**Zone:** From the **Data Management** tab, select the **DNS** tab, click the **Zones** tab -> **DNS View** column -> select the *zone\_name* checkbox, and click the Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the editor, click **Toggle Advanced Mode**, select the **Zone Transfers** tab.
3. In the *Allow zone transfers* section, select one of the following:
  - **None:** Select this to deny all clients for DNS zone transfers. This is selected by default.
  - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this, the appliance allows remote name servers that have the **Allow** permission to send and receive zone transfer data. You can click **Clear** to remove the selected named ACL.
  - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
    - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address of the remote name server. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
    - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
      - **Permission:** Select **Allow** or **Deny** from the drop-down list.
    - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
      - **Permission:** Select **Allow** or **Deny** from the drop-down list.
    - **TSIG Key:** In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
      - **Key name:** Enter a meaningful name for the key, such as a zone name or the name of the remote name server with which the local server authenticates zone transfer requests and replies. This name must match the name of the same TSIG key on other name servers that use it to authenticate zone transfers with the local server.
      - **Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.
      - **Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop down list, and then click **Generate Key Data** to create a new key.

- **DNSone 2.x TSIG Key:** Select this when the other name server is a NIOS appliance running DNS One 2.x code. The appliance automatically populate the value of the key in the **Value** field. The **Permission** column displays **Allow** by default. You cannot change the default permission.
- **Any Address/Network:** Select to allow or deny the local appliance to send zone transfers to any IP address.

After you have added access control entries, you can do the following:

- Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
  - Reorder the list of ACEs using the up and down arrows next to the table.
  - Select an ACE and click the Edit icon to modify the entry.
  - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
4. Optionally, select the **Add allowed IP addresses to also-notify** checkbox to add all IPv4 and IPv6 addresses listed in the "Allow zone transfers to" table to the also-notify statement for each authoritative zone served by a Grid member. When you enable this, all external secondary servers that are not defined for the zone and are allowed zone transfers will receive notifications about zone updates, in addition to name servers assigned to the zone. Infoblox recommends that you do not configure a large number of external secondary servers in stealth mode. To ensure that these secondary servers receive notifications about zone updates, add their addresses to the "Allow zone transfers to" table and grant them the "Allow" permission, and then select this checkbox. Note that the appliance includes only IPv4 and IPv6 addresses. It does not include network addresses, TSIG keys, and denied addresses. When you configure a named ACL, all allowed IPv4 and IPv6 addresses in the named ACL are added to the also-notify statement.
  5. Optionally, you can:
    - Modify an item on the list by selecting it and clicking the Edit icon.
    - Remove an item from the list by selecting it and clicking the Delete icon.
    - Move an item up or down the list. Select it and drag it to its new position, or click the up or down arrow.
  6. Save the configuration and click **Restart** if it appears at the top of the screen.

### Specifying a Zone Transfer Format

The zone transfer format determines the BIND format for a zone transfer. This provides tracking capabilities for single or multiple transfers and their associated servers.

To specify a zone transfer format:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab, click the **Members** tab -> *member* checkbox -> Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Advanced** subtab of the **Zone Transfers** tab to specify the zone transfer format. Select one of the following options from the **Default Zone Transfer Format** drop-down menu:
  - **many-answers** (Secondaries run BIND 8/9): includes as many records as the packet size allows
  - **one-answer** (Secondaries run BIND 4): includes one record per packet
4. To exclude servers, click the Add icon in the **Zone Transfer Format Exceptions** table and enter the IP address of the server in the Addresses field.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

### Configuring Concurrent Zone Transfers

The default number of zone transfers that are allowed is set at the Grid or member level. However, you can override the default value and configure the required concurrent zone transfers. Note that when you increase the number of concurrent zone transfers, there will be an impact on CPU and memory usage.





#### Note

The `tcp-client` value is unconditionally set to 1000 to control the total number of simultaneous TCP connections, which cap the maximum inbound and maximum outbound transfer plus any DNS request made with the TCP. The `tcp-client` value specifies the maximum number of simultaneous DNS clients that can be handled with TCP connections and does not account for UDP connections. The UDP connection accounts for the regular DNS requests and TCP is used only for AXFR and rare DNS requests that don't fit in a UDP connection. You can change the `tcp-client` value by running the `set named_tcp_clients_limit` command.

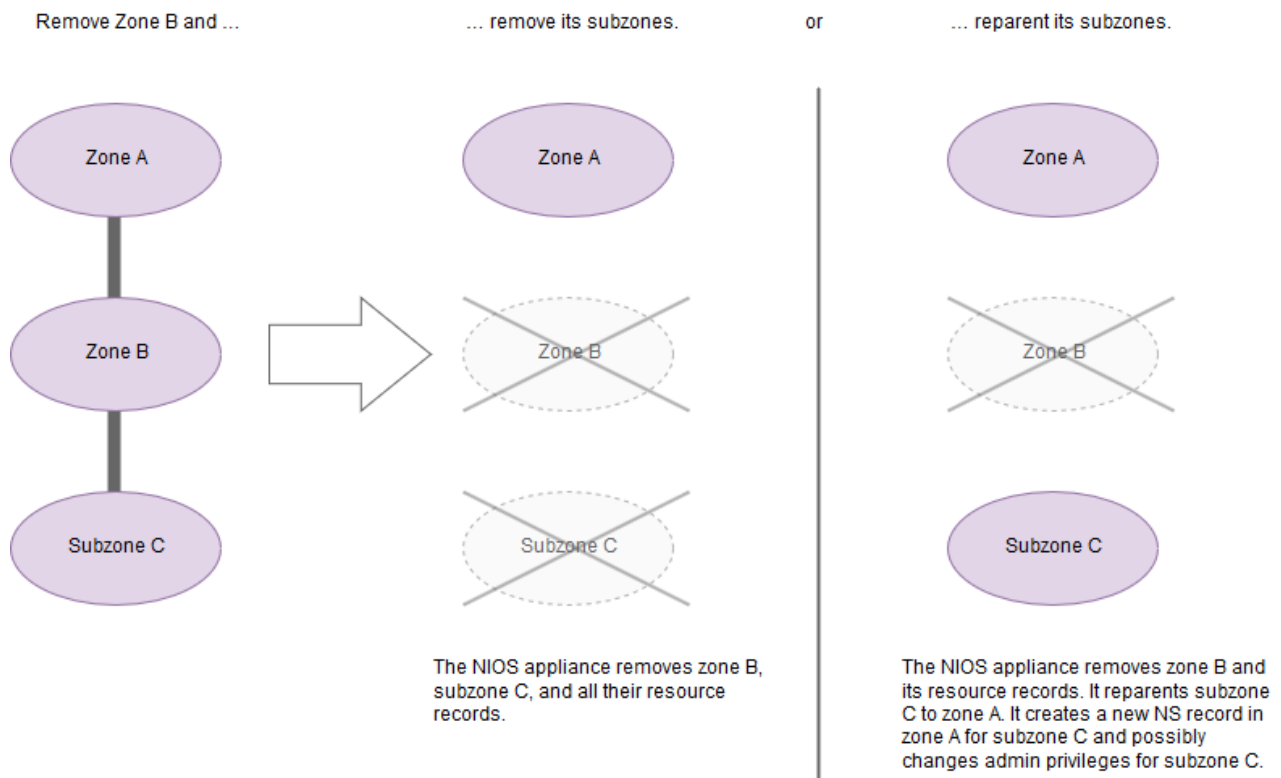
To specify concurrent zone transfers:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Independent appliance:** From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties** -> **Edit**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> **Edit**.
2. In the editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Advanced** subtab of the **General** tab.
4. You can change the zone transfer settings as follows:
  - Maximum inbound concurrent zone transfers: The maximum number of inbound zone transfers that can be performed concurrently. Click **Override** to override the value inherited from the Grid and enter the required value. The default value is 10. Make sure that you specify a value from 10 to 100. Otherwise, the appliance displays an error message. To retain the same value as the Grid, click **Inherit**.
  - Maximum outbound concurrent zone transfers: The maximum number of outbound zone transfers that can be performed concurrently. Click **Override** to override the value inherited from the Grid and enter the required value. The default value is 10. Make sure that you specify a value from 1 to 100. Otherwise, the appliance displays an error message. To retain the same value as the Grid, click **Inherit**.
  - Maximum concurrent inbound zone transfers per remote name server: The maximum number of zone transfers that can be performed concurrently from a given remote name server. This configuration can be done on a per server basis. Click **Override** to override the value inherited from the Grid and enter the required value. The default value is 2. Make sure that you specify a value from 2 to 100. Otherwise, the appliance displays an error message. To retain the same value as the Grid, click **Inherit**.
  - Maximum concurrent SOA queries: The maximum number of concurrent queries a secondary name server sends to the primary server to find out if the zone serial numbers have been changed. Click **Override** to override the value inherited from the Grid and enter the required value. The default value is 20. Make sure that you specify a value from 20 to 1000. Otherwise, the appliance displays an error message. To retain the same value as the Grid, click **Inherit**.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Removing Zones

Depending on the configuration, you may or may not be able to delete or schedule the deletion of a zone and all its contents. Superusers can determine which group of users are allowed to delete or schedule the deletion of a zone and all its contents. For information about how to configure the recursive deletion of zones, see [Configuring Recursive Deletions of Networks and Zones](#).

Note that you must have Read/Write permission to all the subzones and resource records in order to delete a zone. The possible effects of removing or re-parenting are illustrated in the following figure [Removing or Reparenting Subzones](#). The appliance puts all deleted objects in the Recycle Bin, if enabled. You can restore the objects if necessary. When you restore a parent object from the Recycle Bin, all its contents, if any, are re-parented to the restored parent object. For information about the Recycle Bin, see [Using the Recycle Bin](#).  
*Removing or Reparenting Subzones*



If you choose to reparent the subzones, be aware of the following caveats and possible effects of the reparenting:

- You cannot remove a zone and reparent its subzones if at least one of the subzones is a delegated zone. You must first remove any delegated subzones, and then you can remove the zone and reparent its subzones.
- If there are AD (Active Directory) subzones (`_msdcs`, `_sites`, `_tcp`, `_udp`, `domaindnszones`, `foresetdnszones`) and you opt to remove the parent zone only, the NIOS appliance reparents all subzones except the AD subzones, which it removes regardless of the removal option you specify.
- The subzone reparenting option is unavailable when you select multiple zones for removal.
- A record created under a top-level reverse-mapping zone is reparented when its immediate parent zone is created. If that parent zone is deleted, the record is restored to the top-level reverse-mapping zone.

#### Examples:

##### Example 1:

Step 1 - Add `10.in-addr.arpa` under `.` (root zone)

Step 2 - If you add `10.in-addr.arpa`, it is created under `.` (root zone)

Step 3 - if you add `in-addr.arpa`, then `10.in-addr.arpa` is reparented under `in-addr.arpa`

##### Example 2

- Deleting `in-addr.arpa` from the hierarchy might lead to `10.in-addr.arpa` reparenting under `.` (root zone), depending on the **Remove zone only/ Remove all subzones** option you select.
- If `in-addr.arpa` is restored, it is restored under `.` (root) zone with all its resource records.

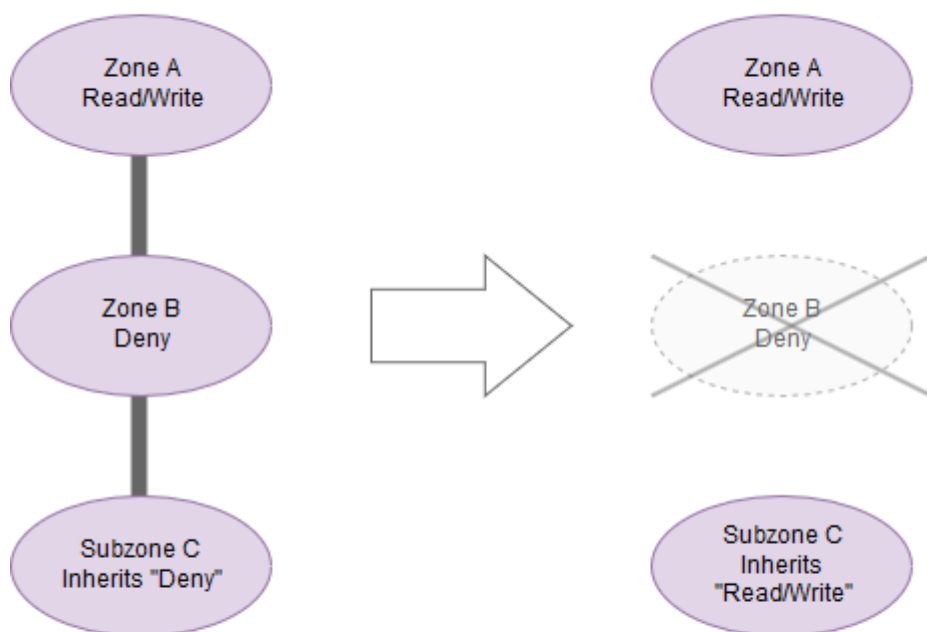
##### Example 3

- Consider `in-addr.arpa` zone having `10.10.in-addr.arpa + 10.0.0.1` (PTR record)
- If you add `10.in-addr.arpa`, then `10.10.in-addr.arpa` is reparented under `10.in-addr.arpa`
- and `10.0.0.1` PTR record is reparented from `in-addr.arpa` to `10.in-addr.arpa`.
- If you delete `10.in-addr.arpa`, then `10.10.in-addr.arpa` is reparented under `in-addr.arpa` (depending on the **Remove zone only/ Remove all subzones** option) and `10.0.0.1` PTR record is deleted along with `10.in-addr.arpa` zone.
- When you remove a zone and reparent its subzones, any subzone that inherited its admin access settings from its previous parent zone (as opposed to having specific access settings for the subzone) now receive their settings from its new parent zone, which might be different. See the following figure *Changed Admin Access Settings after Reparenting Subzones*.

*Changed Admin Access Settings after Reparenting Subzones*

Remove Zone B and ...

... reparent its subzones.



... the admin access settings for subzone C change because the privileges for its new parent zone (zone A) are different from those of its previous parent zone (zone B).

Before you remove zone B, subzone C inherits a "Deny" admin access setting from zone B. After the removal, subzone C inherits "Read/Write" access from its new parent zone, zone A.

Note that if you set a specific "Deny" admin access privilege for subzone C before removing its parent zone (zone B), subzone C retains its specified "Deny" setting.

... the admin access settings for subzone C change because the privileges for its new parent zone (zone A) are different from those of its previous parent zone (zone B).

Before you remove zone B, subzone C inherits a "Deny" admin access setting from zone B. After the removal, subzone C inherits "Read/Write" access from its new parent zone, zone A.

Note that if you set a specific "Deny" admin access privilege for subzone C before removing its parent zone (zone B), subzone C retains its specified "Deny" setting.



#### Note

Instead of removing a zone, you can also disable it. For more information, see [Enabling and Disabling Zones](#).

To remove a zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab.
2. Click the checkbox of the zones you want to delete.
3. Click the Delete icon.
4. Select one of the following. Note that these options appear only if you are allowed to delete zones and all its contents. For information about how to configure this, see [Configuring Recursive Deletions of Networks and Zones](#).
  - **Remove zone only:** Select this to remove the zone and all its content. The appliance reparents all subzones to the parent zone of the zone that you want to remove, except for the automatically created AD (Active Directory) subzones.
  - **Remove all subzones:** Select this to remove the selected zone, all its subzones, and all the resource records of the selected zone and its subzones.
5. Click **Yes**. Grid Manager displays a warning message. Click **Yes** to continue or **No** to cancel the process. Note that this process may take a longer time to complete depending on the size of the data.

You can also schedule the deletion for a later time. Click **Schedule Deletion** and in the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Deletions](#). For information about scheduling recursive deletions of zones, see [Scheduling New IPAM/DHCP Objects and Associated Port Configurations](#).

## Restoring Zone Data

After you import or delete a zone, if you want the original zone back, you can restore it using the Recycle Bin. When you import a zone for the first time, the appliance saves the zone and its resource records as a single object in the Recycle Bin. It keeps the subzones with the zone. See [Restoring Zone Data After a Zone Import Example](#) below.

When you reimport data into a zone, the software saves the zones, its resource records, and the delegated subzones created by the previous import operation in the Recycle Bin. It keeps the subzones (not created during the zone import) with the zone. See [Restoring Zone Data After a Zone Reimport Example](#) below.

If the zone import succeeds, the system adds resource records from the source to the target zone. It also adds delegated subzones for the source subzones. If the zone import fails, the system does not create records and delegated subzones. In either case, you can retrieve the original zone and its subzones from the Recycle Bin as follows:

1. Delete the zone using the steps described in the section [Removing Zones](#).
2. Select **Remove zone only** to remove the zone and its resource records. The NIOS appliance reparents all subzones to the parent zone of the zone that you remove. Do not select **Remove all subzones**. Automatically created AD (Active Directory) subzones are an exception. Even if you select **Remove zone only**, the NIOS appliance still removes AD subzones.
3. In the Finder panel, click **Recycle Bin**.
4. Select the zone you want to restore and click the Restore icon. Click **Yes** in the *Restore Item* dialog box to restore or **No** to cancel the process. Note that restoring a zone may take a longer time to complete depending on the size of the data.  
The zone is restored back to its original state. The resource records are reparented back under it.

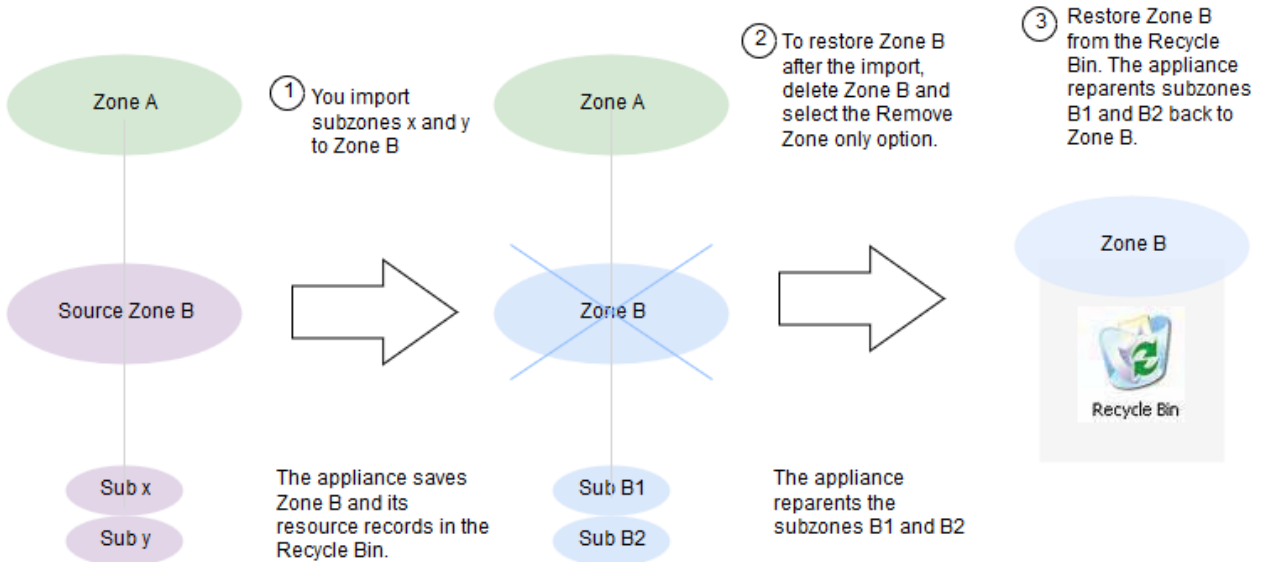
### Restoring Zone Data After a Zone Import Example

In the example shown in the figure *Restoring Zones After a Zone Import* below:

1. Import data from a source zone with subzones Sub x and Sub y into zone B with subzones Sub B1 and Sub B2. The appliance stores zone B and its resource records in the Recycle Bin. To retrieve zone B after the import:
2. Delete subzone B using the **Remove zone only** option.  
The appliance reparents subzones Sub B1 and Sub B2 to the Zone A, which is the zone above Zone B.

- After the import, you can restore zone B from the Recycle Bin. The appliance reparents the subzones Sub B1 and Sub B2 back to zone B.

### Restoring Zones After a Zone Import

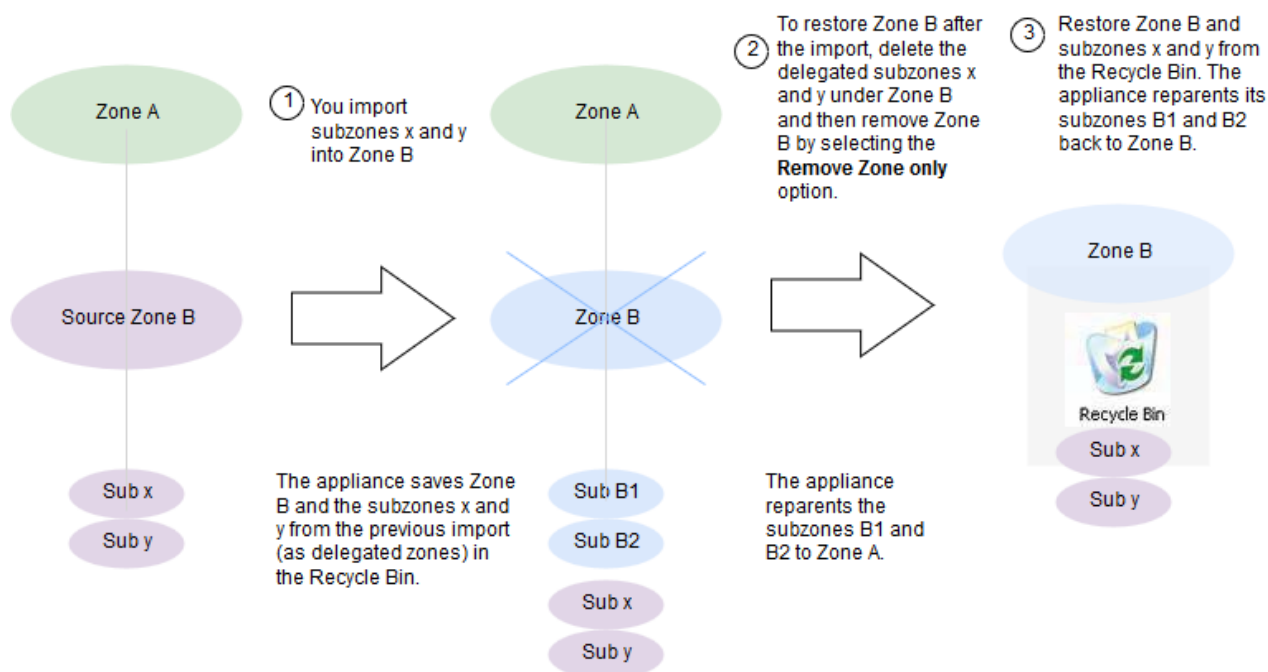


### Restoring Zone Data After a Zone Reimport Example

In the example shown in the figure Restoring Zones After a Zone Reimport below:

- You reimport data from the source zone with subzones Sub x and Sub y into zone B with subzones Sub B1 and Sub B2.  
To retrieve zone B after the import:
- Delete the delegated subzones x and y and then remove subzone B using the **Remove zone only** option. The appliance stores zone B and its resource records and the previously-imported subzones Sub x and Sub y (as delegated subzones) in the Recycle Bin. It reparents subzones Sub B1 and Sub B2 to the zone above zone B (Zone A).
- After the import, you can restore zone B and the subzones Sub x and Sub y from the Recycle Bin. The appliance reparents the subzones Sub B1 and Sub B2 back to zone B.

### Restoring Zones After a Zone Reimport



## Configuring Delegated, Forward, and Stub Zones

In addition to authoritative zones, the NIOS appliance allows you to configure delegated, forward, and stub zones. A delegated zone is a zone managed by (delegated to) another name server who owns the authority for the zone. A forward zone is where queries are sent before being forwarded to other remote name servers. A stub zone contains records that identify the authoritative name servers in another zone. This section covers the following topics:

- [Configuring a Delegation](#)
- [Configuring a Delegation for a Forward-Mapping Zone](#)
- [Configuring a Delegation for a Reverse-Mapping Zone](#)
- [Configuring a Forward Zone](#)
- [Configuring Stub Zones](#)
  - [Creating Stub Zones](#)
  - [Maintaining Stub Zones](#)
  - [Adding Stub Zones](#)
  - [Viewing and Modifying SOA Records](#)
  - [Configuration Example: Configuring a Stub Zone in a Grid](#)

### Configuring a Delegation

Instead of a local name server, remote name servers (which the local server knows) maintain delegated zone data. When the local name server receives a query for a delegated zone, it either responds with the NS record for the delegated zone server (if recursion is disabled on the local server) or it queries the delegated zone server on behalf of the resolver (if recursion is enabled).

For example, there is a remote office with its own name servers, and you want it to manage its own local data. On the name server at the main corporate office, define the remote office zone as delegated, and then specify the remote office name servers as authorities for the zone.

You can delegate a zone to one or more remote name servers, which are typically the authoritative primary and secondary servers for the zone. If recursion is enabled on the local name server, it queries multiple delegated name servers based on their round-trip times. You can also add arpa as a top-level forward-mapping zone and delegate its subzones.

You can also configure TTL settings of auto-generated NS records and glue A and AAAA records for delegated zones in forward-mapping, IPv4 reverse-mapping, and IPv6 reverse-mapping zones. For information, see [Specifying Time To Live Settings](#).

The delegation must exist within an authoritative zone with a Grid primary server.

## Configuring a Delegation for a Forward-Mapping Zone

To create a delegation for a forward-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab.
2. Click the parent zone to open it.  
Grid Manager displays the **Records** and **Subzones** tabs of the zone.
3. From the **Subzones** tab, click the Add icon -> **Zone** -> **Add Delegation**.
4. In the *Add Delegation* wizard, specify the following:
  - **Name:** This field displays a dot followed by the domain name of the current zone. Enter one or more labels before the dot to specify the domain name of the subzone.
  - **DNS View:** This field displays only when there is more than one DNS view in the network view. Displays the DNS view of the current zone.
  - **Comment:** Optionally, enter additional text about the zone.
  - **Disable:** Click this checkbox to temporarily disable this zone. For information, see [Enabling and Disabling Zones](#). Note that disabling a zone may take a longer time to complete depending on the size of the data.
  - **Lock:** Click this checkbox to lock the zone so that you can make changes to it, and also prevent others from making conflicting changes. For information, see [Locking and Unlocking Zones](#).
5. Click **Next** to assign a delegation name server group or define the name servers for the zone. Select one of the following:
  - **Use this nameserver group:** Select this to assign a delegation NS group for the delegated zone. You can select the delegation NS group from the drop-down list.
  - **Use this set of nameservers:** Select this to define name servers for the delegated zone. In the *Name Servers* panel, click the Add icon and specify the following information:
    - **Name:** Enter the name of a remote name server to which you want the local server to redirect queries for zone data. This is a name server that is authoritative for the delegated zone.
    - **Address:** Enter the IP address of the delegated server.

For information about delegation NS group, see [Using Delegation Name Server Groups](#).
6. Save the configuration and click **Restart** if it appears at the top of the screen, or click **Next** to define extensible attributes as described in [Using Extensible Attributes](#). or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).



### Note

The DNS server resolves the FQDN of the delegated name server and does not use the IP address that you specify when assigning the delegated name servers.

## Configuring a Delegation for a Reverse-Mapping Zone

To create a delegation for a reverse-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab.
2. Click the parent zone to open it.  
Grid Manager displays the **Records** and **Subzones** tabs of the zone.
3. From the **Subzones** tab, click the Add icon -> **Zone** -> **Add Delegation**.
4. In the *Add Delegation* wizard, specify the following:
  - **IPv4 Network:** This field displays if you are creating a delegation zone for an IPv4 reverse-mapping zone. Enter the IPv4 address for the address space for which you want to define the reverse-mapping zone and



select a netmask from the Netmask drop-down list. Alternatively, you can specify the address in CIDR format, such as 192/8.

- To use an RFC 2317 prefix, select a netmask value that is between 25 to 31, inclusive. Grid Manager displays the following fields:
    - **RFC2317 Prefix:** Enter a prefix in this field. Prefixes can include alphanumeric characters.
    - **Allow manual creation of PTR records in parent zone:** Select this checkbox to allow users to create labels that correspond to IP addresses in the delegated address space in the parent zone.
    - For information about RFC 2317, see [Specifying an RFC 2317 Prefix](#).
    - **IPv6 Network Prefix:** This field displays if you are creating a delegation zone for an IPv6 reverse-mapping zone. Enter the IPv6 prefix for the address space for which you want to define the reverse-mapping zone and select the prefix length from the drop-down list.
    - **Name:** This field displays a dot followed by the domain name of the current zone. Enter one or more labels before the dot to specify the domain name of the subzone.
    - **DNS View:** This field displays only when there is more than one DNS view in the network view. Select a DNS view from the drop-down list.
    - **Comment:** Optionally, enter additional text about the zone.
    - **Disable:** Select this option to temporarily disable this zone. Note that disabling a zone may take a longer time to complete depending on the size of the data.
    - **Lock:** Select this option to lock the zone so that you can make changes to it and prevent others from making conflicting changes.
5. Click **Next** to assign a delegation name server group or define the name servers for the zone. Select one of the following:
- **Use this name server group:** Select this to assign a delegation NS group for the delegated zone. You can select the delegation NS group from the drop-down list.
  - **Use this set of name servers:** Select this to define name servers for the delegated zone. In the *Name Servers* panel, click the Add icon and specify the following information:
    - **Name:** Enter the name of a remote name server to which you want the local server to redirect queries for zone data. This is a name server that is authoritative for the delegated zone.
    - **Address:** Enter the IP address of the delegated server.
- For information about delegation NS groups, see [Using Delegation Name Server Groups](#).
6. Save the configuration and click **Restart** if it appears at the top of the screen, or click **Next** to define extensible attributes as described in [Using Extensible Attributes](#). or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).



#### Note

The DNS server resolves the FQDN of the delegated name server and does not use the IP address that you specify when assigning the delegated name servers.

## Configuring a Forward Zone

When you want to forward queries for data in a particular zone, define the zone as a forward zone and specify one or more name servers that can resolve queries for the zone. You can also assign one or more external name servers as default forwarders for a forward zone. For example, define a forward zone so that the NIOS appliance forwards queries about a partner's internal site to a name server, which the partner hosts, configured just for other partners to access.

You can override the default forwarders for a forward-mapping zone at a Grid member level and configure custom forwarders. In other words, each Grid member can have its own forwarders for the forward zone. For example: a forward-mapping zone foo.com served by two Grid members M1 and M2 with M1 forwarding queries to 10.1.0.1 and 10.1.0.2 and M2 forwarding queries to 90.3.3.3 and 90.4.4.1. Note that the Grid member uses the default forwarders unless you override them at any level. For more information about domains and zones, see [Configuring Authoritative Zone Properties](#).





#### Note

The use of a forward zone is different from that of a forwarder. (A forwarder is a name server that performs recursive lookups on behalf of the name servers that forward queries to it. For more information, see [Using Forwarders](#).) A NIOS appliance forwards queries to the name server of a forward zone because the name server can resolve queries for the zone. A NIOS appliance forwards queries to a forwarder regardless of zones.

The NIOS appliance automatically generates a name server record in the parent authoritative zone when a subdomain associated with the parent is conditionally forwarded. For example, consider that you configure the following within a single DNS view:

- An authoritative zone *parent.tld* where Grid member is the default primary.
- A subdomain *subdomain.parent.tld*, which is a conditional forwarding zone, and forwards queries to the *ns1.abczone.tld* zone with the IP address *1.2.3.4*.

NIOS automatically creates an *RDATA ns1.abczone.tld* name server record for *subdomain.parent.tld* in the *parent.tld* authoritative zone. You can however disable the generation of name server records and the NIOS appliance deletes all the existing name server records from the parent zone.

Note that a name server can have only one definition for a zone in any given DNS view; a forward zone cannot be configured on a member that already has a zone with the same domain name configured on it in the same DNS view. To configure a forward-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Forward Zone**.
2. In the *Add Forward Zone* wizard, click **Add a forward forward-mapping zone** and click **Next**.
3. Enter the following information, and then click **Next**:
  - **Name:** Enter the domain name of the zone for which you want the NIOS appliance to forward queries.
  - **DNS View:** This field displays only when there is more than one DNS view in the current network view. Select the DNS view of the forward zone.
  - **Comment:** Enter a descriptive comment.
  - **Disable:** Click this checkbox to temporarily disable this zone. Note that disabling a zone may take a longer time to complete depending on the size of the data.
  - **Lock:** Click this checkbox to lock the zone so that you can make changes to it and prevent others from making conflicting changes.
4. Click **Next** to assign a forward/stub server name server group or define the default zone forwarders to which the NIOS appliance forwards queries for the zone. Select one of the following:
  - a. Select **Use this name server group** to assign a forward/stub server NS group for the zone. You can select the forward/stub server NS group from the drop-down list. For information about forward/stub server NS groups, see [Using Forward/Stub Server Name Server Groups](#).
  - b. Select **Use this set of name servers** to specify the default servers for the zone. Click the Add icon and specify the following:
    - **Name:** Enter a domain name of the server to which you want the NIOS appliance to forward queries.
    - **Address:** Enter the IP address of the server to which you want the NIOS appliance to forward queries.
  - c. Select **Disable auto-generation of NS records in parent authoritative zone** to disable generation of name server records in a parent authoritative zone that has a subzone, which is conditionally forwarded. The NIOS appliance will not generate name server records and deletes the existing records from the parent authoritative zone when you select the checkbox. Note that the checkbox is clear, by default, which means that the NIOS appliance automatically generates name server records in a parent authoritative zone.
  - d. Select **Use Forwarders Only** if you want the NIOS appliance to query forwarders only (not root servers) to resolve domain names in the zone.
5. Click **Next** to assign a forwarding member name server group or define Grid members to serve the forward-mapping zone. Select one of the following:

Note that if you do not define any Grid members to serve the forward-mapping zone, then the named.conf file will not contain the configuration of the newly created forward zone. Hence, the Infoblox DNS server will not be authoritative to the forward zone and by default, the Infoblox DNS server will query the root servers to resolve queries for the forward zone.

- a. Select **Use this name server group** to assign a forwarding member NS group for the zone. You can select the forwarding member NS group from the drop-down list. For information about forwarding member NS groups, see [Using Forwarding Member Name Server Groups](#).
- b. Select **Use this set of name servers** to define the Grid members and use the default forwarders or you can override default forwarders and configure custom forwarders. Click the Add icon to select the NIOS appliance on which the forward zone is configured. For an independent deployment, select the local appliance (it is the only choice). If there are multiple Grid members, the *Member Selector* dialog box is displayed. Select the required member by clicking the member name.

The following is displayed for each Grid member:

- **Name:** Displays the name of the Grid member.
- **IPv4 Address:** Displays the IPv4 address of the Grid member.
- **IPv6 Address:** Displays the IPv6 address of the Grid member.
- **Override Default Forwarders:** Displays **Yes** when you override default forwarders. Otherwise, this field displays **No**.
- **Custom Forwarders:** Displays the IP address of the custom forwarders. Otherwise, this field is blank.

Note to skip the following two steps if you want to use the default forwarders.

6. Select a member and click the Edit icon.
7. In the *Edit Per-Member Forwarders* editor, select the **Override Default Forwarders** checkbox to override the default forwarders. The Default Zone Forwarders table becomes available only after you select the **Override Default Forwarders** checkbox. Click the Add icon to specify the servers to which the NIOS appliance forwards queries for the zone:
  - **Name:** Enter a domain name for the server to which you want the NIOS appliance to forward queries for the specified domain name.
  - **Address:** Enter the IP address of the server to which you want the NIOS appliance to forward queries.
  - Select **Use Forwarders Only** if you want the NIOS appliance to query forwarders only (not root servers) to resolve domain names in the zone.
  - Save the configuration. After successfully saving the configuration, the **Override Default Forwarders** column displays **Yes** and the **Custom Forwarders** column displays the IP address of the forwarders. To configure forwarders for multiple members, repeat the steps for each Grid member.
8. Save the configuration, or click **Next** to continue to the next step where you define extensible attributes as described in [Using Extensible Attributes](#), and then optionally proceed to the next step where you define admin permissions as defined in [About Administrative Permissions](#).  
or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).
9. Click **Restart** if it appears at the top of the screen.

To configure a forward IPv4 reverse-mapping zone:

1. From the **Data Management** tab, select the **Zones** tab, expand the Toolbar and click **Add -> Zone -> Add Forward Zone**.
2. In the *Add Forward Zone* wizard, click **Add a forward IPv4 reverse-mapping zone** and click **Next**.
3. Enter the following information, and then click **Next**:
  - a. **IPv4 Network:** Enter the IPv4 address for the address space for which you want to define the reverse-mapping zone and select a netmask from the **Netmask** drop-down list. Alternatively, you can specify the address in CIDR format, such as 192/8.  
To use an RFC 2317 prefix, select a netmask value that is between 25 to 31, inclusive. Grid Manager displays the **RFC 2317 Prefix** field. Enter a prefix in the text field. Prefixes can be alphanumeric characters. For information, see [Specifying an RFC 2317 Prefix](#).  
or  
**Name:** Enter the domain name of the reverse-mapping zone.
  - b. **DNS View:** This field displays only when there is more than one DNS view in the network view. Select a DNS view from the drop-down list.
  - c. **Comment:** Optionally, enter additional information about the zone.
  - d. **Disable:** Click this checkbox to temporarily disable this zone. Note that disabling a zone may take a longer time to complete depending on the size of the data.
  - e. **Lock:** Click this checkbox to lock the zone so that you can make changes to it, and also prevent others from making conflicting changes.

4. Click **Next** to assign a forward/stub server name server group or define the default zone forwarders to which the NIOS appliance forwards queries for the zone. Select one of the following:
  - Select **Use this name server group** to assign a forward/stub server NS group for the zone. You can select the forward/stub server NS group from the drop-down list. For information about forward/stub NS groups, see [Using Forward/Stub Server Name Server Groups](#).
  - Select **Use this set of name servers** to specify the default servers for the zone. Click the Add icon and specify the following:
    - **Name:** Enter a domain name for the server to which you want the NIOS appliance to forward queries.
    - **Address:** Enter the IP address of the server to which you want the NIOS appliance to forward queries.
    - Select **Use Forwarders Only** if you want the NIOS appliance to query forwarders only (not root servers) to resolve domain names in the zone.
5. Click **Next** to assign a forwarding member name server group or define Grid members to serve the forward-mapping zone. Select one of the following:
  - Select **Use this name server group** to assign a forwarding member NS group for the zone. You can select the forwarding member NS group from the drop-down list. For information about forwarding member NS groups, see [Using Forwarding Member Name Server Groups](#).
  - Select **Use this set of name servers** to define the Grid members and use the default forwarders or you can override default forwarders and configure custom forwarders. Click the Add icon to select the NIOS appliance on which the forward zone is configured. For an independent deployment, select the local appliance (it is the only choice). If there are multiple Grid members, the *Member Selector* dialog box is displayed. Select the required member by clicking the member name.
 

The following is displayed for each Grid member:

    - **Name:** Displays the name of the Grid member.
    - **IPv4 Address:** Displays the IPv4 address of the Grid member.
    - **IPv6 Address:** Displays the IPv6 address of the Grid member.
    - **Override Default Forwarders:** Displays **Yes** when you override default forwarders. Otherwise, this field displays **No**.
    - **Custom Forwarders:** Displays the IP address of the custom forwarders. Otherwise, this field is blank.

Note to skip the following two steps if you want to use the default forwarders.
6. Select a member and click the Edit icon.
7. In the *Edit Per-Member Forwarders* editor, select the **Override Default Forwarders** checkbox to override the default forwarders. The Default Zone Forwarders table becomes available only after you select the **Override Default Forwarders** checkbox. Click the Add icon to specify the servers to which the NIOS appliance forwards queries for the zone:
  - a. **Name:** Enter a domain name for the server to which you want the NIOS appliance to forward queries for the specified domain name.
  - b. **Address:** Enter the IP address of the server to which you want the NIOS appliance to forward queries.
  - c. Select **Use Forwarders Only** if you want the NIOS appliance to query forwarders only (not root servers) to resolve domain names in the zone.
  - d. Save the configuration. After successfully saving the configuration, the **Override Default Forwarders** column displays **Yes** and the **Custom Forwarders** column displays the IP address of the forwarders. To configure forwarders for multiple members, repeat the steps for each Grid member.
8. Save the configuration, or click **Next** to continue to the next step where you define extensible attributes as described in [Using Extensible Attributes](#).  
or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).
9. Click **Restart** if it appears at the top of the screen.

To configure a forward IPv6 reverse-mapping zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Forward Zone**.
2. In the *Add Forward Zone* wizard, click **Add a forward IPv6 reverse-mapping zone** and click **Next**.
3. Enter the following zone information:

- **IPv6 Network Address:** Enter the 128-bit IPv6 address for the address space for which you want to define the reverse-mapping zone. When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2006:0000:0000:0123:4567:89ab:0000:cdef can be shortened to 2006::123:4567:89ab:0:cdef. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The NIOS appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered. Choose the network prefix that defines the IPv6 network address space.
  - or
  - **Name:** Enter the domain name of the reverse-mapping zone.
  - **DNS View:** This field displays only when there is more than one DNS view in the network view. Select a DNS view from the drop-down list.
  - **Comment:** Enter a descriptive comment about the zone.
  - **Disable:** Click this checkbox to temporarily disable this zone. Note that disabling a zone may take a longer time to complete depending on the size of the data.
  - **Lock:** Click this checkbox to lock the zone so that you can make changes to it, and also prevent others making conflicting changes.
4. Click **Next** to assign a forward/stub server name server group or define the default zone forwarders to which the NIOS appliance forwards queries for the zone. Select one of the following:
    - Select **Use this name server group** to assign a forward/stub server NS group for the zone. You can select the forward/stub server NS group from the drop-down list. For information about forward/stub NS groups, see [Using Forward/Stub Server Name Server Groups](#).
    - Select **Use this set of name servers** to specify the default servers for the zone. Click the Add icon and specify the following:
      - **Name:** Enter a domain name for the server to which you want the NIOS appliance to forward queries.
      - **Address:** Enter the IP address of the server to which you want the NIOS appliance to forward queries.
      - Select **Use Forwarders Only** if you want the NIOS appliance to query forwarders only (not root servers) to resolve domain names in the zone.
  5. Click **Next** to assign a forwarding member name server group or define Grid members to serve the forward-mapping zone. Select one of the following:
    - Select **Use this name server group** to assign a forwarding member NS group for the zone. You can select the forwarding member NS group from the drop-down list. For information about forwarding member NS groups, see [Using Forwarding Member Name Server Groups](#).
    - Select **Use this set of name servers** to define the Grid members and use the default forwarders or you can override default forwarders and configure custom forwarders. Click the Add icon to select the NIOS appliance on which the forward zone is configured. For an independent deployment, select the local appliance (it is the only choice). If there are multiple Grid members, the *Member Selector* dialog box is displayed. Select the required member by clicking the member name.
 

The following is displayed for each Grid member:

      - **Name:** Displays the name of the Grid member.
      - **IPv4 Address:** Displays the IPv4 address of the Grid member.
      - **IPv6 Address:** Displays the IPv6 address of the Grid member.
      - **Override Default Forwarders:** Displays **Yes** when you override default forwarders. Otherwise, this field displays **No**.
      - **Custom Forwarders:** Displays the IP address of the custom forwarders. Otherwise, this field is blank.

## Configuring Stub Zones

A stub zone contains records that identify the authoritative name servers in the zone. It does not contain resource records for resolving IP addresses to hosts in the zone. Instead, it contains the following records:

- SOA (Start of Authority) record of the zone
- NS (name server) records at the apex of the stub zone
- A (Address) records that map the name servers to their IP addresses

Stub zones, like secondary zones, obtain their records from other name servers. Their records are read only; therefore, administrators do not manually add, remove, or modify the records.

Stub zone records are also periodically refreshed, just like secondary zone records. However, secondary name servers contain a complete copy of the zone data on the primary server. Therefore, zone transfers from a primary server to a secondary server, or between secondary servers, can increase CPU usage and consume excessive bandwidth. A name server hosting a stub zone maintains a much smaller set of records; therefore, updates are less CPU intensive and consume less bandwidth.

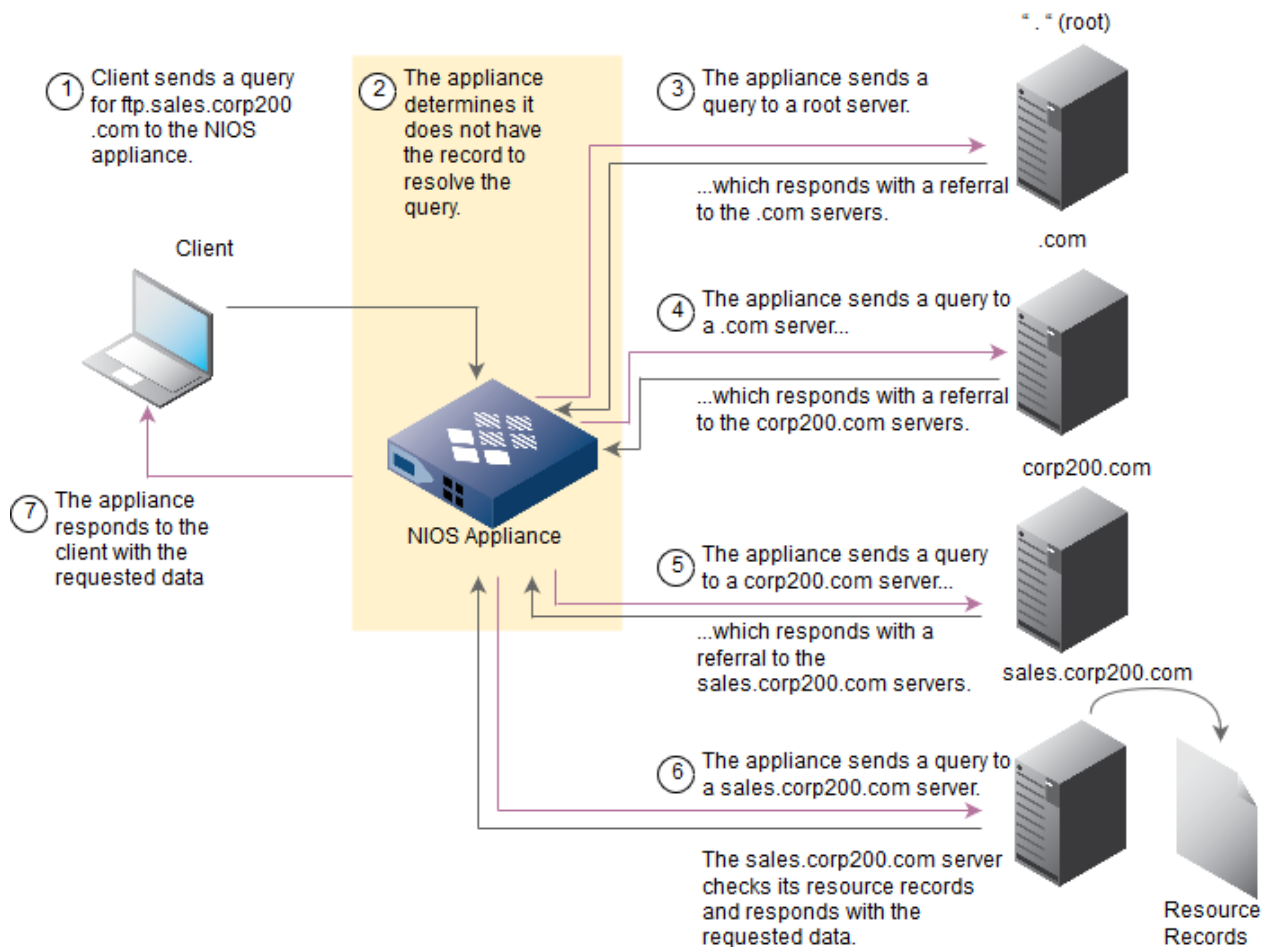
When a name server hosting a stub zone receives a query for a domain name that it determines is in the stub zone, the name server uses the records in the stub zone to locate the correct name server to query, eliminating the need to query the root server.

The figures *Processing a Query without a Stub Zone* and *Processing a Query with a Stub Zone* below illustrate how the NIOS appliance resolves a query for a domain name for which it is not authoritative. The following figure illustrates how the appliance resolves a query when it does not have a stub zone.

The figure *Processing a Query with a Stub Zone* illustrates how the appliance resolves the query with a stub zone.

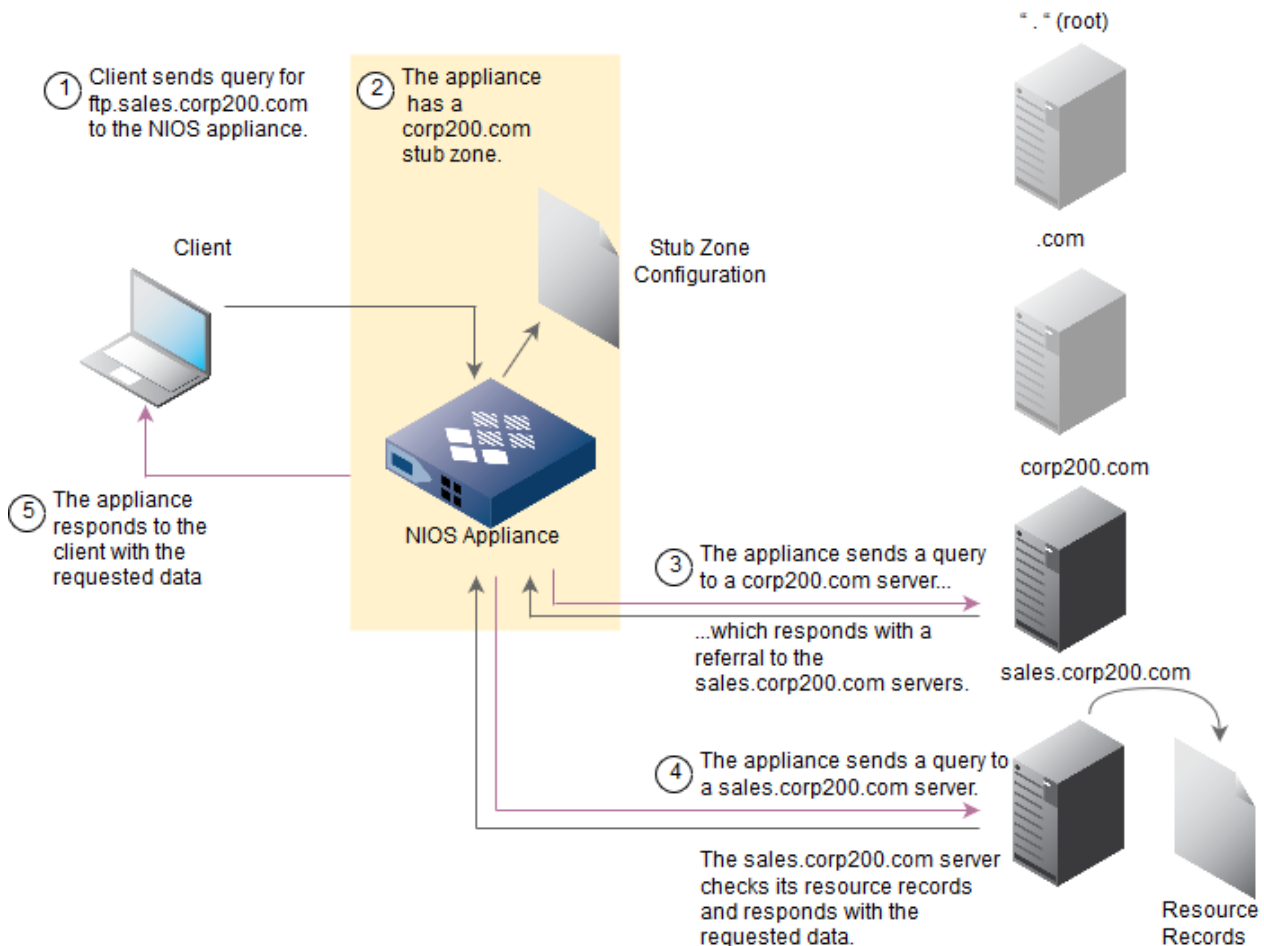
In the figure *Processing a Query without a Stub Zone*, a client sends a query for ftp.sales.corp200.com to the NIOS appliance. When the appliance receives the request from the client, it checks if it has the data to resolve the query. If the appliance does not have the data, it tries to locate the authoritative name server for the requested domain name. It sends nonrecursive queries to a root name server and to the closest known name servers until it learns the correct authoritative name server to query.

*Processing a Query without a Stub Zone*



In the following figure Processing a Query with a Stub Zone, when the NIOS appliance receives the request for the domain name in corp200.com, it determines it does not have the resource records to resolve the query. It does, however, have a list of the authoritative name servers in the stub zone, corp200.com. The appliance then sends a query directly to the name server in corp200.com.

*Processing a Query with a Stub Zone*




Stub zones facilitate name resolution and alleviate name server traffic in your network. For example, the client in the previous examples is in corpxyz.com. The corpxyz.com and corp200.com zones are partners, and send all their communications through a VPN tunnel, as shown in the figure Stub Zone Configuration above. The firewall protecting corpxyz.com is configured to send all messages for the 10.2.2.0/24 network through the VPN tunnel. Infoblox\_A hosts the stub zone for corp200.com. Therefore, when the host in corpxyz.com sends a query for ftp.sales.corp200.com, Infoblox\_A obtains the IP address of Infoblox\_B (10.2.2.7) from its stub zone records and sends the query to the firewall protecting corpxyz.com.

Because the destination of the query is in the 10.2.2.0/24 network, the firewall (configured to encrypt all traffic to the network) sends the request through a VPN tunnel to Infoblox\_B. Infoblox\_B resolves the query and sends back the response through the VPN tunnel. All name server traffic went through the VPN tunnel to the internal servers, bypassing the root servers and external name servers.

*Stub Zone Configuration*





In parent-child zone configurations, using stub zones also eases the administration of name servers in both zones. For example, as shown in figure Stub Zone Configuration above, sales.corp200.com is a child zone of corp200.com. On the corp200.com name servers, you can create either a delegated zone or a stub zone for sales.corp200.com.

When you create a delegated zone, you must first specify the name servers in the delegated zone and manually maintain information about these name servers. For example, if the administrator in sales.corp200.com changes the IP address of a name server or adds a new name server, the sales.corp200.com administrator must inform the corp200.com administrator to make the corresponding changes in the delegated zone records.

If, instead, you create a stub zone for sales.corp200.com, you set up the stub zone records once, and updates are then done automatically. The name servers in corp200.com that are hosting a stub zone for sales.corp200.com automatically obtain updates of the authoritative name servers in the child zone.

In addition, a name server that hosts a stub zone can cache the responses it receives. Therefore, when it receives a request for the same resource record, it can respond without querying another name server.

## Creating Stub Zones

When you create a stub zone on the NIOS appliance, you specify the following:

- The Grid member that is hosting the stub zone.  
You can specify multiple appliances if you want the stub zones on multiple name servers. If you do, the appliances store identical records about the stub zone. You can also specify a stub member NS group for the zone. For information on specifying a stub member NS group, see [Using Stub Member Name Server Groups](#).
- The IP address of the primary server(s) that the NIOS appliance can query in the stub zone.

The primary server can be a Grid member or an external primary server. If you specify multiple primary servers, the appliance queries the primary servers, starting with the first server on the list. You can also specify a forward/stub server NS group for the zone. For information on specifying a forward/stub server NS group, see [Using Forward/Stub Server Name Server Groups](#).

The primary server and the name server hosting the stub zone can belong to the same Grid, as long as the authoritative zone and the stub zone are in different DNS views. You cannot configure one zone as both authoritative and stub in the same view.

After you create a stub zone, the NIOS appliance does the following:

1. It sends a query to the primary server for the SOA (Start of Authority) record of the stub zone. The primary server returns the SOA record.
2. Then, it sends a query for the NS (name server) records in the zone.  
The primary server returns the NS records and the A (address) records of the name servers. (These A records are also called glue records.)  
If the primary server is a NIOS appliance, you might have to manually create the A record and add it to the stub zone. A NIOS appliance that is the primary server for a zone always creates an NS record, but does not always create an A record.

- The appliance automatically creates an A record when its host name belongs to the name space of the zone. For example, if the zone is corpxyz.com and the primary server host name is server1.corpxyz.com, the appliance automatically creates the NS and A records and sends these records when it is queried by the stub zone name server.
- The appliance does not automatically create an A record when its host name is in a name space that is different from the zone. For example, if the zone is corp200.com and the primary server host name is server1.corpxyz.com, then the appliance creates the NS record only and sends it when it is queried by the stub zone name server. In this case, you must manually create the A record.

## Maintaining Stub Zones

The NIOS appliance maintains the stub zone records and updates them based on the values in the SOA record as follows:

- The refresh interval indicates when the appliance sends a discrete query to the primary name server for the stub zone. The appliance learns about any changes in the stub zone and updates the NS and A records in the stub zone accordingly.
- If the update fails, the retry interval indicates when the appliance resends a discrete query.
- If the query continues to fail, the expiry value indicates when the appliance stops using the zone data.

## Adding Stub Zones

To add a stub zone, you must identify the Infoblox appliance that hosts the stub zone, and provide the IP address of the primary server.

You can also add stub zones for Microsoft servers that are managed by Grid members. For information, see [Managing Microsoft Windows Servers](#).

You can configure a stub zone for forward mapping or reverse mapping zones.

To add a forward-mapping stub zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Stub Zone**.
2. In the *Add Stub Zone* wizard, click **Add a stub forward-mapping zone** and click **Next**.
3. Specify the following, and then click **Next**:
  - **Name**: Enter the name for the stub zone.
  - **Comment**: Enter a useful comment, such as the admin to contact for the stub zone.
  - **Disable**: Click this checkbox to temporarily disable this zone. Note that disabling a zone may take a longer time to complete depending on the size of the data.
  - **Lock**: Click this checkbox to lock the zone so that you can make changes to it, and also prevent others from making conflicting changes.
4. Click **Next** to define primary servers in the stub zone. You can specify a forward/stub server NS group or define the servers individually. Select one of the following:
  - Select **Use this name server group** to assign a forward/stub server NS group for the stub zone. You can select the forward/stub server NS group from the drop-down list. For information about forward/stub server NS groups, see [Using Forward/Stub Server Name Server Groups](#).
  - Select **Use this set of name servers** to define primary servers for the stub zone. Click the Add icon and enter the **Name** and **IP Address** of the primary server in the stub zone.
 

If the primary server is a Grid member, you must enter the host name and IP address of the Grid member. The NIOS appliance does not validate these entries. Therefore, if you change the IP address of a Grid member listed here, you must update the Grid member information in this list as well.

You can specify multiple primary servers for redundancy. If the primary server is a NIOS appliance, the appliance must have the Minimal Response feature disabled so it can propagate the data to the stub server. For information about the Minimal Response feature, see [Specifying Minimal Responses](#).

Optionally, click the **Don't use forwarders to resolve queries in subzones** checkbox to indicate that the name servers hosting the stub zone must not use forwarders to resolve queries for domain names in the stub zone or in its subzones.



5. Click **Next** to specify a stub member NS group or define the name servers individually to serve the forward-mapping stub zone. Select one of the following:
  - Select **Use this name server group** to assign a stub member NS group for the zone. You can select the stub member NS group from the drop-down list. For information about stub member NS group, see [Using Stub Member Name Server Groups](#).
  - Select **Use this set of name servers** to define the servers individually. Click the Add icon and select one of the following:
    - **Add Infoblox Member:** Select this and select the Grid member that hosts the stub zone.
    - **Add Microsoft Server:** Select this and select the Microsoft server that hosts the stub zone. The following is displayed for each name server:
      - **Name:** Displays the name of the name server.
      - **IPv4 Address:** Displays the IPv4 address of the name server.
      - **IPv6 Address:** Displays the IPv6 address of the name server.
6. Click **Next** to continue to the next step where you define extensible attributes as described in [Using Extensible Attributes](#).
7. Save the configuration and click **Restart** if it appears at the top of the screen or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

You can define two types of reverse-mapping stub zones, one for IPv4 addresses and one for IPv6 addresses. To configure an IPv4 reverse-mapping stub zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Stub Zone**.
2. In the *Add Stub Zone* wizard, click **Add a stub IPv4 reverse-mapping zone** and click **Next**.
3. Specify the following:
  - **IPv4 Network:** Enter the IPv4 address for the address space for which you want to define the reverse-mapping zone and select a netmask from the **Netmask** drop-down list. Alternatively, you can specify the address in CIDR format, such as 192/8.
  - To use an RFC 2317 prefix, select a netmask value that is between 25 to 31, inclusive. Grid Manager displays the **RFC 2317 Prefix** field. Enter a prefix in the text field. Prefixes can be alphanumeric characters. For information, see [Specifying an RFC 2317 Prefix](#).
  - or
  - **Name:** Enter the domain name of the reverse-mapping zone.
  - **DNS View:** This field displays only when there is more than one DNS view in the network view. Select a DNS view from the drop-down list.
  - **Comment:** Optionally, enter additional information about the zone.
  - **Disable:** Click this checkbox to temporarily disable this zone. Note that disabling a zone may take a longer time to complete depending on the size of the data.
  - **Lock:** Click this checkbox to lock the zone so that you can make changes to it, and also prevent others from making conflicting changes.
4. Click **Next** to define primary servers in the stub zone. You can specify a forward/stub server NS group or define the servers individually. Select one of the following:
  - Select **Use this name server group** to assign a forward/stub server NS group for the stub zone. You can select the forward/stub server NS group from the drop-down list. For information about forward/stub server NS group, see [Using Forward/Stub Server Name Server Groups](#).
  - Select **Use this set of name servers** to define primary servers for the stub zone. Click the Add icon and enter the **Name** and **IP Address** of the primary server in the stub zone.  
If the primary server is a Grid member, you must enter the host name and IP address of the Grid member. The NIOS appliance does not validate these entries. Therefore, if you change the IP address of a Grid member listed here, you must update the Grid member information in this list as well.  
You can specify multiple primary servers for redundancy. If the primary server is a NIOS appliance, the appliance must have the Minimal Response feature disabled so it can propagate the data to the stub server. For information about the Minimal Response feature, see [Specifying Minimal Responses](#).  
Optionally, click the **Don't use forwarders to resolve queries in subzones** checkbox to indicate that the name servers hosting the stub zone should not forward queries that end with the domain name of the stub zone to any configured forwarders.

5. Click **Next** to specify a stub member NS group or define the name servers individually to serve the reverse-mapping stub zone. Select one of the following:
  - Select **Use this name server group** to assign a stub member NS group for the zone. You can select the forward/stub server NS group from the drop-down list. For information about stub member NS group, see [Using Stub Member Name Server Groups](#).
  - Select **Use this set of name servers** to define the servers individually. Click the Add icon and select one of the following:
    - **Add Infoblox Member:** Select this and select the Grid member that hosts the stub zone.
    - **Add Microsoft Server:** Select this and select the Microsoft server that hosts the stub zone. The following is displayed for each name server:
      - **Name:** Displays the name of the name server.
      - **IPv4 Address:** Displays the IPv4 address of the name server.
      - **IPv6 Address:** Displays the IPv6 address of the name server.
6. Click **Next** to continue to the next step where you define extensible attributes as described in [Using Extensible Attributes](#).
7. Save the configuration and click **Restart** if it appears at the top of the screen or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

To configure an IPv6 reverse-mapping stub zone:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Stub Zone**.
2. In the *Add Stub Zone* wizard, click **Add a stub IPv6 reverse-mapping zone** and click **Next**.
3. Specify the following:
  - **IPv6 Network Prefix and Prefix Length:** Enter the 128-bit IPv6 address for the address space for which you want to define the reverse-mapping zone. When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2006:0000:0000:0123:4567:89ab:0000:cdef can be shortened to 2006::123:4567:89ab:0:cdef. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The NIOS appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered. You can enter a slash and prefix length in the IPv6 Network Prefix field or you can choose a value from the Prefix Length drop-down list.  
or
  - **Name:** Enter the domain name of the reverse-mapping zone.
  - **DNS View:** This field displays only when there is more than one DNS view in the current network view. Select a DNS view from the drop-down list.
  - **Comment:** Enter a descriptive comment about the zone.
  - **Disable:** Click this checkbox to temporarily disable this zone. Note that disabling a zone may take a longer time to complete depending on the size of the data.
  - **Lock:** Click this checkbox to lock the zone so that you can make changes to it and prevent others from making conflicting changes.
4. Click **Next** to define primary servers in the stub zone. You can specify a forward/stub server NS group or define the servers individually. Select one of the following:
  - Select **Use this name server group** to assign a forward/stub server NS group for the stub zone. You can select the forward/stub server NS group from the drop-down list. For information about forward/stub server NS group, see [Using Forward/Stub Server Name Server Groups](#).
  - Select **Use this set of name servers** to define primary servers for the stub zone. Click the Add icon and enter the **Name** and **IP Address** of the primary server in the stub zone.  
If the primary server is a Grid member, you must enter the host name and IP address of the Grid member. The NIOS appliance does not validate these entries. Therefore, if you change the IP address of a Grid member listed here, you must update the Grid member information in this list as well.  
You can specify multiple primary servers for redundancy. If the primary server is a NIOS appliance, the appliance must have the Minimal Response feature disabled so it can propagate the data to the stub server. For information about the Minimal Response feature, see [Specifying Minimal Responses](#).  
Optionally, click the **Don't use forwarders to resolve queries in subzones** checkbox to indicate that the

- name servers hosting the stub zone should not forward queries that end with the domain name of the stub zone to any configured forwarders.
5. Click **Next** to specify a stub member NS group or define the name servers individually to serve the reverse-mapping stub zone. Select one of the following:
    - Select **Use this name server group** to assign a stub member NS group for the zone. You can select the stub member NS group from the drop-down list. For information about stub member NS group, see [Using Stub Member Name Server Groups](#).
    - Select **Use this set of name servers** to define the servers individually. Click the Add icon and select one of the following:
      - **Add Infoblox Member**: Select this and select the Grid member that hosts the stub zone.
      - **Add Microsoft Server**: Select this and select the Microsoft server that hosts the stub zone. The following is displayed for each name server:
        - **Name**: Displays the name of the name server.
        - **IPv4 Address**: Displays the IPv4 address of the name server.
        - **IPv6 Address**: Displays the IPv6 address of the name server.
  6. Click **Next** to continue to the next step where you define extensible attributes as described in [Using Extensible Attributes](#).
  7. Save the configuration and click **Restart** if it appears at the top of the screen or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Viewing and Modifying SOA Records

The timer values in the SOA record determine when the zone data is updated. The MNAME field and the RNAME field of the SOA record display the FQDN of the primary server and the administrative email address respectively. You can view these default values and override them when necessary. For a zone that has multiple primary servers, Grid Manager displays all configured primaries for the zone. You can click **Override** to override the Grid-level settings. If the primary server is a Microsoft server however, the **Override** option does not appear. You can only change certain values in the SOA record.

To view and modify zone SOA record values:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox, and then click the Edit icon.
2. In the *Authoritative Zone* or *Stub Zone* editor, click the **Settings** tab to view the following values. You can also click **Override** to modify some of the values.
  - **Serial number**: The current serial number for the primary server. This number is automatically increased when changes are made to the zone or its record. The serial number plays a key role in determining when and whether zone data is updated. You can change the serial number only if the primary server of the zone is a Grid member. When the zone has multiple primary servers, each primary can have its own serial number. In this case, the serial number displayed here is always that of the Grid Master, which will also appear in the primary name server list if it is one of the primaries for the zone.  
Note that if you change the serial number of the Grid Master, serial numbers for all primaries will be changed to the same number. A warning is displayed when you try to decrement the serial number.
  - **Refresh**: This interval tells secondary servers how often to send a message to the primary server for a zone to check that their data is current, and retrieve fresh data if it is not. The default is three hours.
  - **Retry**: This interval tells the secondary server how long to wait before attempting to recontact the primary server after a connection failure between the two occurs. The default is one hour.
  - **Expire**: If the secondary fails to contact the primary for the specified interval, the secondary stops giving out answers about the zone because the zone data is too old to be useful. The default is 30 days.
  - **Default TTL**: Specifies how long name servers can cache the data. The default is eight hours.
  - **Negative-caching TTL (Time to Live)**: Specifies how long name servers can cache negative responses. The default is 15 minutes.
  - **Primary name server (for SOA MNAME field)**: If the primary name server of a zone is a Grid member, the MNAME is inherited from its corresponding member, and you can change the name of the primary name server that is published in the MNAME field of the SOA record. This field accepts names in native character sets. If the zone has multiple primary name servers, a list of all primaries is displayed in this section. Each primary has its own serial number and the number can be different among them. Note that

- the serial numbers for these primaries are read-only and you cannot modify them. If you change the serial number of the Grid Master, serial numbers for all primaries will be changed to the same number.
- **Email Address (for SOA RNAME field):** If the primary name server of a zone is a Grid member, you can enter an administrator email address to the SOA record to help people determine who to contact about this zone. The appliance supports IDN for the host name of the Email address. For example, you can create admin@инфоблокс.рф but not админ@инфоблокс.рф.com.
  - **Don't use forwarders to resolve queries in subzones:** Select this option to disable the use of forwarders to resolve queries for data in subzones.
3. Save the configuration and click **Restart** if it appears at the top of the screen. To schedule this task, click the **Schedule** icon at the top of the wizard. In the **Schedule Change** panel, click **Later**, and then specify a date, time, and time zone. The **Schedule** icon is green when there is a pending scheduled task. You can reschedule the task if you have the applicable permissions.

### Configuration Example: Configuring a Stub Zone in a Grid

This example illustrates how to configure a stub zone and assign it to a Grid member. You configure a Grid, corpxyz, with a single Grid Master and Grid member. The Grid member, member1.corpxyz.com, is the primary name server for the corpxyz.com zone in the internal view. The Grid Master, gm-corpxyz.com, hosts the stub zone for corpxyz.com in the external view. Thus, when the Grid Master receives a query for the corpxyz.com zone, it sends it directly to member1.corpxyz.com, the primary name server for the zone.

In this example, you configure the following:

1. Turn off minimal responses on member1.corpxyz.com, the primary name server for the corpxyz.com zone. See [Disable Minimal Responses](#) below.
2. Create the internal and external views. See [Create the Views](#) below.
3. Create the corpxyz.com authoritative zone and stub zone. See [Create the Zones](#) below.

#### Disable Minimal Responses

After you create the Grid, turn off minimal responses for member1.corpxyz.com. Disabling minimal responses ensures that member1.corpxyz.com propagates the required data to the server hosting the stub zone.

1. From the **Data Management** tab, select the **DNS** tab, click **Members** -> **member1.corpxyz.com** checkbox -> Edit icon.
2. In the *Member DNS Configuration* editor, click the **General** -> **Basic** tab.
3. Clear the **Return minimal responses** checkbox.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

#### Create the Views

Create the internal and external views. To create each view:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add DNS View**.
2. In the *Add DNS View* wizard, enter the name of the view. In this example, enter either **External** or **Internal**.
3. Click **Save & New** and create the other DNS view.

#### Create the Zones

Create the corpxyz.com zone in the internal view and assign member1.corpxyz.com as the Grid primary server:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add** -> **Zone** -> **Add Auth Zone**.
2. In the *Forward Authoritative Zone* wizard, do the following:
  - Select **Add an authoritative forward-mapping zone** and click **Next**.
  - Enter the zone name, **corpxyz.com** and select the **Internal** DNS view. Click **Next**.
  - Select **Use this set of name servers** and select **member1.corpxyz.com** as the Grid primary server.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

After you create the zone, you can view the NS and A records which were automatically created.

Create the stub zone, corpxyz.com, in the external view, assign gm-corpxyz.com as the stub member and member1.corpxyz.com as the stub primary server.

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Zone -> Add Stub Zone**.
2. In the *Stub Zone* wizard, do the following:
  - Select **Add a stub forward-mapping zone** and click **Next**.
  - Enter the name of the stub zone, **corpxyz.com** and select the **External** DNS view. Click **Next**.
  - In the *Master Name Servers* panel, click the Add icon and enter the following for the primary name server, and then click **Next**:
    - **Name:** member1.corpxyz.com
    - **Address:** 10.35.0.222
  - In the *Name Servers* panel, click the Add icon and select **gm-corpxyz.com**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

After you create the stub zone, the server hosting the stub zone, gm-corpxyz.com, sends queries to the primary server, member1.corpxyz.com, for the SOA and NS records. member1.corpxyz.com then returns its NS records and A (address) records.

## Viewing Zones

To list zones, navigate to the **Data Management** tab -> **DNS** tab -> *Zones* panel. If there is more than one DNS view in the Grid, this panel lists the DNS views. Select a DNS view to list its zones. (For information, see [Listing DNS Views.](#))

- Click Toggle flat view to display a flat list of all the zones in the view.
- Click Toggle hierarchical view to display only the apex zones.

In the hierarchical view, you can see one entry for the host that represents the entire host object. In a host record, there can be multiple DNS resource records (A, PTR, CNAME) and some DHCP data (fixed addresses) as well. In the flat view, each of the DNS resource records in the host are listed separately.

For example, the host called server1.infoblox.com contains 2 A records and an ALIAS (which is a host naming convention for CNAME records). If you view the infoblox.com zone using the hierarchical view option, you will see one entry host for server1.infoblox.com. In the flat view, you will see three records (one for each IP address/A record, and one host Alias for the CNAME). In the flat view, you cannot delete one piece of the host record. You can edit the host record and you can remove information. Deleting host records deletes the entire host record only.

This panel displays the following information for each zone, by default:

- **Name:** The domain name of the zone.
- **MS Sync Server:** When a zone is served by multiple Microsoft servers, this column shows which Microsoft server is actually performing the synchronization of that zone with the Grid.
- **MS Zone Sync:** Displays **Yes** if you have enabled zone synchronization, and displays **No** when the zone synchronization is disabled. This column appears only when you have the Microsoft license installed.
- **Grid Primary Server:** The primary name server configured for an authoritative zone in the DNS view.
- **Type:** The zone type. Possible values are Authoritative, Forward, Stub and Delegation.
- **Multi-master Zone:** Indicates whether this zone has multiple primary name servers.
- **Comment:** Comments that were entered for the zone.
- **Site:** Values that were entered for this pre-defined attribute.

You can also display the following columns:

- **Locked:** Displays Yes when a zone is locked by an admin, and displays No when the zone is unlocked.
- **Function:** Indicates whether the zone is a forward-mapping, or an IPv4 or IPv6 reverse-mapping zone.
- **ZSK rollover date:** Displays the date when the ZSK is due for next rollover. The appliance performs a rollover automatically at this time.
- **KSK rollover date:** Displays the date when the KSK is due for next rollover. The appliance performs a rollover automatically at this time if you have enabled automatic KSK rollover. For more information, see [Configuring Automatic KSK Rollovers and Notifications](#). You must perform the rollover manually, if you have disabled this option.
- **Disabled:** This field displays **Yes** if the zone is disabled. Otherwise, this field displays **No**. You can double click a row and select the checkbox in this column to disable the zone. Grid Manager displays a warning message when you disable the zone. Click **Yes** to confirm or **No** to cancel.
- **Signed:** This field displays Yes if the zone is a DNSSEC-signed zone. Otherwise, this field displays No. You can do the following:

- List the resource records and subzones of a DNS zone.
  - Click a DNS zone name.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).
- Edit the properties of a DNS zone.
  - Click the checkbox beside a DNS zone, and then click the Edit icon.
- Delete a DNS zone.
  - Click the checkbox beside a DNS zone, and then click the Delete icon. Grid Manager displays a warning message. Click **Yes** to continue or **No** to cancel the process. Note that this process may take a longer time to complete depending on the size of the data.
- Export the list of DNS zones to a .csv file.
  - Click the Export icon.
- Print the list of DNS zones.
  - Click the Print icon.

## Using Name Server Groups

A name server (NS) group is a collection of one or more DNS servers. Grouping a commonly used set of DNS servers together simplifies zone creation by enabling you to specify a single NS group instead of specifying multiple name servers individually.

You can configure the following types of name server groups:

- Authoritative name server group
- Delegation name server group
- Forwarding member name server group
- Stub member name server group
- Forward/Stub server name server group

### Using Authoritative Name Server Groups

An authoritative name server group is a collection of one or more primary DNS servers and secondary DNS servers. After you create an authoritative NS group, you can then assign it to serve authoritative forward-mapping and reverse-mapping zones. When you assign an authoritative NS group to an authoritative zone, Grid Manager automatically generates an NS record, a glue A or AAAA record, and a PTR record for each name server available in the NS group. But if the zone is disabled, Grid Manager does not generate these records.

Grid Manager generates authoritative A, AAAA, and PTR records when:

- A Grid member is added to a Grid, whose host name belongs to the name space of the authoritative zone and vice versa.
- An external name server is assigned to a zone, whose host name belongs to the name space of the authoritative zone.



#### Note

Grid Manager does not generate an NS record when the DNS service for the member is disabled.

The performance of the following functions significantly improve when you assign an NS group to a zone instead of specifying multiple name servers individually:

- Starting and Stopping the DNS service.
- Reparenting the zones after removing or restoring a zone.
- Modifying the zone data.





#### Note

Only superusers can create and manage name server groups

## Adding Authoritative Name Server Groups

To add an authoritative name server group:

1. From the **Data Management** -> **DNS** tab, do one of the following:
  - Click the **Name Server Groups** tab -> Add icon -> **Group** -> **Authoritative**.
  - From the Toolbar, click the Add icon -> **Group** -> **Authoritative**.
2. In the *Add Name Server Group* wizard, do the following:
  - **Name:** Type a name that provides a meaningful reference for this set of servers.
  - **Name Servers:** Click the Add icon and select one of the following options for every server that you are adding to the NS group:
    - **Grid Primary:** Choose this option to select a primary name server or multiple primary servers for the zone. See [Specifying Grid Primary Servers](#).
    - **Grid Secondary:** Choose this option to select a Grid member as a secondary server for the zone. See [Adding Grid Secondaries](#).
    - **External Primary:** Choose this option if the appliance is in a Grid and you want to specify a primary server outside the Grid ("external" to the Grid). See [Specifying External Primary Servers](#).
    - **External Secondary:** Choose this option if the appliance is in a Grid and you want to specify a secondary server outside the Grid ("external" to the Grid), or if the appliance is deployed independently from a Grid. See [Specifying External Secondaries](#).
  - **Default NS Group:** Select this to specify this authoritative name server group as the default.
  - **Comment:** Optionally, enter additional information about the authoritative NS group.
3. Save the configuration and click **Restart** if it appears at the top of the screen, or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).

A newly created authoritative name server group appears in the **Name Server Groups** tab. For information about viewing the name server groups, see [Viewing Name Server Groups](#). You can then associate it with forward-mapping and reverse-mapping authoritative zones.

## Applying Authoritative Name Server Groups

In Grid Manager, you can assign an authoritative name server group to an authoritative zone when you first create it using the *Add Authoritative Zone* wizard and when you edit an existing authoritative zone using the *Authoritative Zone* editor. For information on creating an authoritative zone using the wizard, see [Configuring Authoritative Zones](#). The panels used to assign a name server group to a zone are the same in the wizard or and editor. The only difference is the way you access it. The following procedure describes how to specify an authoritative NS group when editing a forward-mapping zone:

1. From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *zone* checkbox, and then click the Edit icon.
2. In the *Authoritative Zone* editor, click **Name Servers**.
3. Select **Use this name server group**, and then select the authoritative NS group from the drop-down list.



#### Note

If you apply a name server group to at least one zone or specify it as the default group, you cannot rename or remove it. To rename or remove a group, you must first disassociate it from all zones and unassign it as the default group.

## Using Delegation Name Server Groups

A delegation name server group is a collection of external name servers for delegated zones. Adding a set of external name servers to a single delegation name server group can significantly reduce configuration efforts. While configuring delegated zones, you can specify a single delegation NS group instead of configuring multiple name servers individually.

After you create a delegation NS group, you can then assign it to serve only delegated zones. When you assign a delegation NS group to a delegated zone, Grid Manager automatically generates an NS record, a glue A or AAAA record, and a PTR record for each name server available in the delegation NS group. But if the zone is disabled, Grid Manager does not generate these records.

Note the following while adding a delegation NS group:

- You cannot add a delegation NS group if a NS group with the same name already exists and vice versa.
- If a Grid has any Microsoft servers configured, delegation NS groups are not allowed and vice versa.
- You cannot delete a delegation NS group if it is assigned to a zone.
- The DNS server resolves the FQDN of the delegated name server and does not use the IP address that you specify when assigning the delegated name servers.

## Adding Delegation Name Server Groups

To add a delegation name server group:

1. From the **Data Management** -> **DNS** tab, do one of the following:
  - Click the **Name Server Groups** tab -> Add icon -> **Delegation**.
  - From the Toolbar, click the Add icon -> **Group** -> **Delegation**.
2. In the *Add Delegation Name Server Group* wizard, complete the following:
  - **Name:** Enter the name of the delegation NS group.
  - **Comment:** Optionally, enter additional information about the delegation NS group.
3. Click **Next** to define the external name servers for the delegation NS group.
4. In the *Name Server* panel, click the Add icon and specify the following for every server that you are adding to the delegation NS group:
  - **Name:** Enter the name of the delegated name server.
  - **Address:** Enter the IP address of the delegated name server.
5. Save the configuration and click **Restart** if it appears at the top of the screen, or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).

A newly created delegation NS group appears in the **Name Server Groups** tab. For information about viewing the name server groups, see [Viewing Name Server Groups](#). You can then associate it with delegated zones.



### Note

You will not be able to add a delegated name server group if DNS synchronization is enabled on any Microsoft server configured in NIOS.

## Applying Delegation Name Server Groups

You can assign a delegation NS group to a delegated zone when you first create it using the *Add Delegation* wizard and when you edit an existing delegated zone using the *Delegation Zone* editor. For information about creating a delegated zone using the wizard, see [Configuring a Delegation](#).

Complete the following to assign a delegation NS group to a delegation zone:

1. From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *delegation zone* checkbox, and then click the Edit icon.
2. In the *Delegation Zone* editor, click **Delegated Name Servers**.
3. Select **Use this name server group**, and then select the delegation NS group from the drop-down list.

## Using Forwarding Member Name Server Groups

A forwarding member NS (Name Server) group is a collection of one or more forwarding name servers. Grouping a set of forwarding name servers together reduces the configuration efforts. When you configure a forwarding zone, you can specify a single forwarding member NS group instead of specifying multiple forwarding name servers individually. After you create a forwarding member NS group, you can assign it to a forward forward-mapping zone and a forward reverse-mapping zone.

Note the following while adding a forwarding member NS group:

- Only superusers can add, modify, and delete a forwarding member NS group.



- You cannot add a forwarding member NS group if a NS group with the same name already exists and vice versa.
- You cannot delete a forwarding member NS group if it is assigned to a zone.

## Adding Forwarding Member Name Server Groups

To add a forwarding member NS group:

1. From the **Data Management** -> **DNS** tab, do one of the following:
  - Click the **Name Server Groups** tab -> Add icon -> **Forwarding Member**.
  - From the Toolbar, click the Add icon -> **Group** -> **Forwarding Member**.
2. In the *Forwarding Member Name Server Group* wizard, complete the following:
  - **Name:** Enter the name of the forwarding member NS group.
  - **Name Servers:** In this section, you can add the Grid members to the forwarding member NS group and use the default forwarders or you can override default forwarders and configure custom forwarders. Click the Add icon and select a Grid member from the *MemberSelector* dialog box. The following is displayed in the table for each member:
    - **Name:** The name of the Grid member.
    - **IPv4 Address:** The IPv4 address of the Grid member.
    - **IPv6 Address:** The IPv6 address of the Grid member.
    - **Override Default Forwarders:** Displays **Yes** when you override default forwarders. Otherwise, this field displays **No**.
    - **Custom Forwarders:** Displays the IP address of the custom forwarders. Otherwise, this field is blank.
  - **Comment:** Optionally, enter additional information about the forwarding member NS group.
3. Save the configuration or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).

A newly created forwarding member NS group appears in the **Name Server Groups** tab. For information about viewing the name server groups, see [Viewing Name Server Groups](#). You can then associate it with forward forward-mapping zones and forward reverse-mapping zones.

## Applying Forwarding Member Name Server Groups

You can assign a forwarding member NS group to a forward zone when you first create it using the *Add Forward Zone* wizard and when you edit an existing forward zone using the *Forward Zone* editor. For information about creating a forward zone using the wizard, see [Configuring a Forward Zone](#).

Complete the following to assign a forwarding member NS group to a forward zone:

1. From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *forward zone* checkbox, and then click the Edit icon.
2. In the *Forward Zone* editor, click **Name Servers**.
3. Select **Use this name server group**, and then select the forwarding member NS group from the drop-down list.

## Using Stub Member Name Server Groups

A stub member NS (Name Server) Group is a collection of one or more Grid members for stub zones. When you configure a stub zone, you can specify a single stub member NS group instead of specifying multiple Grid members individually, thus reducing the configuration efforts. After you create a stub member NS group, you can assign it to forward-mapping stub zones and reverse-mapping stub zones.

Note the following while adding a stub member NS group:

- Only superusers can add, modify, and delete a stub member NS group.
- You cannot add a stub member NS group if a NS group with the same name already exists and vice versa.
- You cannot delete a stub member NS group if it is assigned to a zone.

## Adding Stub Member Name Server Groups

To add a stub member NS group:

1. From the **Data Management** -> **DNS** tab, do one of the following:

- Click the **Name Server Groups** tab -> Add icon -> **Stub Member**.
  - From the Toolbar, click the Add icon -> **Group** -> **Stub Member**.
2. In the *Stub Member Name Server Group* wizard, complete the following:
    - **Name:** Enter the name of the stub member NS group.
    - **Name Servers:** In this section, you can add the Grid members to the stub member NS group. Click the Add icon and select a Grid member from the *MemberSelector* dialog box. The following is displayed in the table for each member:
      - **Name:** The name of the Grid member.
      - **IPv4 Address:** The IPv4 address of the Grid member.
      - **IPv6 Address:** The IPv6 address of the Grid member.
    - **Comment:** Optionally, enter additional information about the stub member NS group.
  3. Save the configuration or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).

A newly created stub member NS group appears in the **Name Server Groups** tab. For information about viewing the name server groups, see [Viewing Name Server Groups](#). You can then associate it with stub forward-mapping zones and stub reverse-mapping zones.

### Applying Stub Member Name Server Groups

You can assign a stub member NS group to a stub zone when you first create it using the *Add Stub Zone* wizard and when you edit an existing stub zone using the *Stub Zone* editor. For information about creating a stub zone using the wizard, see [Configuring Stub Zones](#).

Complete the following to assign a stub member NS group to a stub zone:

1. From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *stub zone* checkbox, and then click the Edit icon.
2. In the *Stub Zone* editor, click **Name Servers**.
3. Select **Use this name server group**, and then select the stub member NS group from the drop-down list.

### Using Forward/Stub Server Name Server Groups

A forward/stub server NS (Name Server) group is a collection of one or more external name servers. Grouping a set of external name servers together to a forward/stub server NS group reduces the configuration efforts. You can assign a single forward/stub server NS group as default forwarders for a forward zone or as primary servers for a stub zone instead of specifying multiple name servers individually.

Note the following while adding a forward/stub server NS group:

- Only super-users can add, modify, and delete a forward/stub server NS group.
- You cannot add a forward/stub server NS group if a NS group with the same name already exists and vice versa.
- You cannot delete a forward/stub server NS group if it is assigned to a zone.

### Adding Forward/Stub Server Name Server Groups

To add a forward/stub server NS group:

1. From the **Data Management** -> **DNS** tab, do one of the following:
  - Click the **Name Server Groups** tab -> Add icon -> **Forward/Stub Server**.
  - From the Toolbar, click the Add icon -> **Group** -> **Forward/Stub Server**.
2. In the *Forward/Stub Server Name Server Group* wizard, complete the following:
  - **Name:** Enter a name that provides a meaningful reference for this set of external name servers for the forward and stub zones.
  - **Name Servers:** Click the Add icon and specify the following for every external name server that you are adding to the forward/stub server NS group:
    - **Name:** The name of the external name server.
    - **IP Address:** The IP address of the external name server.
  - **Comment:** Optionally, enter additional information about the forward/stub server NS group.
3. Save the configuration or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).

A newly created forward/stub server NS group appears in the **Name Server Groups** tab. For information about viewing the name server groups, see [Viewing Name Server Groups](#). You can then associate it with forward zones as default forwarders and as primary servers for stub zones.

### Applying Forward/Stub Server Name Server Groups

You can assign a forward/stub server NS group to a forward zone as default forwarders and as primary servers to a stub zone when you first create it and when you edit an existing forward zone and stub zone. For information about creating a forward zone using the wizard, see [Configuring a Forward Zone](#). For information about creating a stub zone using the wizard, see [Configuring Stub Zones](#).

Complete the following to assign a forward/stub server NS group to a forward zone:

1. From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *forward zone* checkbox, and then click the Edit icon.
2. In the *Forward Zone* editor, click **Forwarders**.
3. Select **Use this name server group**, and then select the forward/stub server NS group from the drop-down list.  
Complete the following to assign a forward/stub server NS group to a stub zone:
4. From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *stub zone* checkbox, and then click the Edit icon.
5. In the *Stub Zone* editor, click **Masters**.
6. Select **Use this name server group**, and then select the forward/stub server NS group from the drop-down list.

### Viewing Name Server Groups

You can view the configured authoritative, delegation, forwarding member, stub member, and forward/stub server name server groups by navigating to the **Data Management** tab -> **DNS** tab -> **Name Server Groups** tab.

The panel displays the following information about each name server group:

- **Name:** The name of the name server group.
- **Type:** The name server group type. Possible values are **Authoritative**, **Delegation**, **Forwarding Member**, **Stub Member**, and **Forward/Stub Server**.
- **Comment:** Comments that were entered for the name server group.
- **Site:** Values that were entered for this pre-defined attribute. You can do the following:
  - Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
  - Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
  - Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).
- Edit the properties of a name server group.
  - Select the checkbox beside a name server group, and then click the Edit icon.
- Delete a name server group.
  - Select the checkbox beside a name server group, and then click the Delete icon. Note that you cannot delete a delegation name server group that is assigned to a zone.
- Export the list of Grid members to a .csv file.
  - Click the Export icon.
- Print the list of Grid members.
  - Click the Print icon.

## Configuring DNS Resource Records

If you need to add a large number of A and PTR records, you can have the NIOS appliance add them as a group and automatically assign host names based on a range of IP addresses and the host name format you specify. Such a group of records is called a *bulk host*, which the appliance manages and displays as a single bulk host record.

This section provides general information about Infoblox host records and DNS resource records. The topics in this section include:

- [Adding Bulk Hosts](#)
- [Specifying Bulk Host Name Formats](#)
- [Managing Resource Records](#)
- [Configuring Shared Record Groups](#)

## Adding Bulk Hosts

To add a bulk host:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Host -> Add Bulk Host**.
2. In the *Add Bulk Host* wizard, complete the following fields:
  - **Prefix:** If Grid Manager displays a zone name, enter a prefix (or series of characters) to insert at the beginning of each host name. The displayed zone name can either be the last selected zone or the zone from which you are adding the bulk host record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and enter a prefix for the bulk host record. You can enter any printable character that complies with the zone host name policy or you can also leave this blank.  
The sum of the bulk host prefix length and suffix length must not exceed 63 characters. When you enter a prefix, the NIOS appliance computes the maximum length of the bulk host suffix to verify that the total prefix and suffix length does not exceed 63 characters. If it does, the appliance displays an error message indicating the number of characters that you must remove to make a valid prefix.
    - **DNS View:** Displays the DNS view of the zone to which the bulk host records belong.
    - **Host Name Policy:** Displays the host name policy of the selected DNS zone.  
From NIOS 8.6.3 onwards, you have to manually clear the existing host name policy in the table and enter the new one.
    - **Name Format:** To override the default four-octet suffix format or the format set at the Grid level, and specify a different format, click **Override** and select a host name format from the **Name Formats** drop-down menu.
    - The *Name Formats* drop-down menu lists the formats **Four Octets**, **Three Octets**, **Two Octets**, and **One Octet** along with any other bulk host name formats that you have defined.
    - **Starting IP Address:** Enter the first IP address in the range of addresses for the group.
    - **End IP Address:** Enter the last IP address in the range of addresses for the group.
    - **Comment:** Optionally, enter additional information for this record.
    - **Automatically Add Reverse Mapping:** Click to have the appliance automatically create a PTR record for each IP address within the bulk host range.
    - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes for the bulk host record. For information, see [Using Extensible Attributes](#).
4. Save the configuration and click **Restart** if it appears at the top of the screen.  
To modify or delete a bulk host, see [Modifying, Disabling, and Deleting Host and Resource Records](#).

### Example 1 - Responding to DNS AXFR Queries

This example shows the responses the bulk host `foo/1.2.3.10/1.2.3.20` returns to DNS AXFR (Full Zone Transfers) queries.

If the bulk host uses the template `-$3-$4`, the query returns:

```
foo-3-10.test.com foo-3-11.test.com
.....
foo-3-20.test.com
```

If the bulk host uses the template `##2-##3-##4`, the query returns:

```
foo-002-003-010.test.com foo-002-003-011.test.com
.....
foo-002-003-020.test.com
```

## Example 2 - Importing Zones with Bulk Hosts

When you import zones with bulk hosts, the system selects the most specific match.

The following example can possibly match three octet, two octet, and one octet formats; however, the system selects the most specific four octet default format.

The query:

```
foo-1-2-3-4 IN A 1.2.3.4
foo-1-2-3-5 IN A 1.2.3.5
```

Results in the match:

```
foo/1.2.3.4/1.2.3.5(Four Octets)
Not in any of the following:
foo-1/1.2.3.4/1.2.3.5(Three Octets) foo-1-2/1.2.3.4/1.2.3.5(Two Octets)
foo-1-2-3/1.2.3.4/1.2.3.5(One Octet)
```

## Specifying Bulk Host Name Formats

Bulk host name formats provide a flexible way to define bulk host names. You create multiple bulk host formats at the Grid level. Either select from the default bulk host formats or create your own. You can specify a different format for each bulk host. When you assign a bulk host name format to a bulk host in a zone, the system applies the zone's host name policy to it.

A bulk host name consists of a prefix, a suffix, and the name of the domain to which the host belongs. The prefix can contain any printable character that complies with the zone host name policy. It can also be blank. The suffix is derived from an IP address in the bulk host IP address range. The appliance also supports IDNs for bulk host names. You can use IDNs or their punycode representations while creating bulk hosts.

The following table summarizes how the appliance displays bulk host names that contain IDNs:

Input	NIOS Displays...	NIOS DNS Domain (Punycode in the GUI)	Conversion Guidelines
hello	hello	hello	No conversion
привітання	привітання	xn--80adk5aaihr3f9e	IDN to punycode
xn--80adk5aaihr3f9e	xn--80adk5aaihr3f9e	xn--80adk5aaihr3f9e	No conversion
\xyz format	\xyz format	\xyz format	No conversion

The suffix format is a string of ASCII characters that uses \$ (unpadded) or # (zero-padded) followed by 1,2,3,4 to refer to the first, second, third, or fourth IP address octet; it uses \$1,\$2,\$3,\$4 or #1,#2,#3,#4. \$2 refers to the second unpadded octet and #4 refers to the fourth zero-padded octet. For example:

The prefix of a bulk host = *info* IP address = *10.19.32.133* Domain name = *infoblox.com*.  
 If you specify the default four-octet format *\$1\$2-\$3-\$4*, the bulk host name is *info-10-19-23-133.infoblox.com*.

If you specify a custom name format such as `#1#2*#3*#4`, the bulk host name is `info*010*019*023*133.infoblox.com`.

### Before Defining Bulk Host Name Formats

Before you specify a bulk host name format, ensure that it complies with the following rules:

- The NIOS appliance uses `<prefix>-xx-xx-xx-xx` for bulk hosts. Ensure that the bulk host name does not conflict with CNAMEs, DNAMEs, or host name aliases.
- When you add a bulk host, if you enable the **Automatically add reverse mapping** option and there is a CNAME record in the corresponding reverse zone that conflicts with a PTR record generated by the bulk host, the bulk host insertion fails and an error message appears. For example, if there is a CNAME with the alias `15` in a reverse zone `1.168.192.in-addr.arpa` and if you add a bulk host `foo/192.168.1.10/192.168.1.20` with the **Automatically add reverse mapping** option selected, the insertion fails and an error message appears because both the bulk host and the CNAME generate a record `15.1.168.192.in-addr.arpa` in the reverse zone.
- You cannot create or change a bulk host if a zone is locked by another user. If you select a different template for the Grid, it changes each record associated with the bulk host.
- You can define bulk host name formats only at the Grid level and override them at the bulk host level; not at the zone or bulk host object level.
- When you upgrade to NIOS 4.3r3 or earlier releases, the system migrates existing bulk hosts as follows:
  - If you did not customize the bulk host IP format, there is no action required. All migrated bulk hosts continue to use the Grid-level default four-octet format `$1$2-$3-$4`.
  - If you customized the bulk host IP format, the system creates a new template called *Migrated Default* template. All migrated bulk hosts override the Grid default template and use the *Migrated Default* template.



#### Note

The NIOS appliance considers two bulk hosts that have the same prefix, start address, and end address as duplicate hosts; even if they use different bulk host formats.

### Bulk Host Name Format Rules

The below table describes the rules that you should follow when you create bulk host name formats. It also provides examples of valid and invalid formats for each rule.

#### *Bulk Host Name Format Rules and Examples*

Rule	Example
The suffix format cannot have more than four octets.	<code>\$4-\$5</code> is invalid.
The octets must be in order.	<code>-\$2-\$3-\$4</code> is valid but <code>-\$3-\$2-\$4</code> is invalid.
Do not skip octets.	<code>-\$2-\$3-\$4</code> is valid but <code>-\$2-\$4</code> is invalid.
Do not use a combination of both the <b>\$</b> and <b>#</b> symbols together as octet references; use only one of them.	<code>-\$2-#3-\$4</code> is invalid.
The suffix format must contain at least the fourth octet. You must define at least one <code>-\$4</code> or <code>#4</code> .	<code>\$4</code> is valid but <code>\$3</code> is invalid.
If the suffix format uses \$ references, it cannot be preceded by a digit. You must add a non-digit prefix to each \$ or # reference.	<code>-\$2-\$3-\$4</code>

The \ character is the designated escape character for the \$, # and \ characters. You cannot use the \$ or # symbols as separators unless you prefix them with an escape character \.	For the IP address <b>10.19.32.133</b> , the format <b>#-#1-#2-#3-#4</b> expands to #-010-019-032-133.
The bulk host name format must comply with its zone host name policy.	You cannot insert a bulk host name format <b>-?-#4</b> in a zone that uses Allow Underscore as host name policy because the policy does not allow you to use the ? character in the host name.
The bulk host name must comply with the maximum label length.	The sum of the bulk host name prefix and suffix cannot be greater than 63 characters. When you enter a suffix format, the NIOS appliance determines the length of the longest bulk host defined, and checks that the sum of the bulk host prefix and suffix length does not exceed 63 characters; if it does, an error message appears.
The bulk host name cannot result in an FQDN with more than 255 characters.	
The NIOS appliance computes the maximum length of the bulk host suffix by expanding the bulk host IP format using 255.255.255.255.	For the format string <b>-\$1-\$2-\$3-\$4</b> , the maximum length of the suffix is <b>-255-255-255-255</b> ; that is, 16 characters. Therefore, the maximum length of the host prefix is 47 characters.
The bulk host name must not be the same as a CNAME/DNAME.	If there is a CNAME record with alias <b>foo-003-015</b> , you cannot insert a bulk host <b>foo/1.2.3.10/1.2.3.20</b> using template <b>#3#4</b> because <b>foo-003-015</b> is also one of the synthetic host names in the bulk host.
Each host name in the bulk host must be unique.	You cannot insert a bulk host <b>foo/1.2.3.10/1.2.4.20</b> using the template <b>\$4</b> because the system resolves the host name <b>foo-10</b> to both <b>1.2.3.10</b> and <b>1.2.4.10</b> . To ensure that the bulk host name is unique, use the template <b>\$3-\$4</b> .
You cannot insert a bulk host that violates the uniqueness of two bulk hosts that have the same prefix and use the same name format.	If there is a bulk host <b>foo/1.2.3.10/1.2.4.20</b> using the template <b>-\$3-\$4</b> , you cannot insert another bulk host <b>foo/1.3.4.10/1.3.5.20</b> using the same template because the system resolves host name <b>foo-4-15</b> to both <b>1.2.4.15</b> and <b>1.3.4.15</b> . Instead, use the template <b>-\$2-\$3-\$4</b> to ensure that the two bulk hosts are unique.

The appliance provides four predefined formats. You can define additional formats or change the default format at the Grid level only. To define new bulk host name formats:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. Select the Host Naming tab of the Grid DNS Properties editor.  
The Bulk Host Name Formats table displays four predefined name suffix formats. The following examples show the host name that each format generates for the zone test.com:  
Four Octets: \$1-\$2-\$3-\$4 (Default)—Generates foo-192-168-1-15.test.com. Three Octets: \$2\$3-\$4—Generates foo-168-1-15.test.com  
Two Octets: **-\$3-\$4**—Generates foo-1-15.test.com One Octet: **-\$4**—Generates foo-15.test.com  
For the IP address 10.100.0.10, the format **-\$1\$2-\$3-\$4** generates the host name suffix 10-100-0-10 . The format **#1#2-#3-#4** generates the host name suffix -010-100-000-010.
3. Click **Add** to enter the name and format of a new bulk host name format.
4. Optionally, click the Default column of a format and select **Default** to make it the Grid default.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Managing Resource Records

DNS resource records provide information about objects and hosts. DNS servers use these records to respond to queries for hosts and objects. The appliance supports IDNs for all DNS resource records. For information about IDNs, see [Support for Internationalized Domain Names](#). Note that the appliance does not decode the IDN of a resource



record to punycode. In other words, a record that contains a domain name in punycode is displayed in punycode and a record that contains an IDN is displayed in its native characters.

The following sections define the types of DNS resource records you can manage and the operations you can perform:

- [Managing ALIAS Records](#)
- [Managing NS Records](#)
- [Managing AAAA Records](#)
- [Managing PTR Records](#)
- [Managing MX Records](#)
- [Managing SRV Records](#)
- [Managing TXT Records](#)
- [Managing TLSA Records](#)
- [Managing CAA Records](#)
- [Managing CNAME Records](#)
- [Managing DNAME Records](#)
- [Managing NAPTR Records](#)
- [Managing LBDN Records](#)
- [Managing Unknown Records](#)
- [Viewing Resource Records](#)
- [Modifying, Disabling, and Deleting Host and Resource Records](#)

#### Managing A Records

An A (address) record is a DNS resource record that maps a domain name to an IPv4 address. To define a specific name-to-address mapping, you can add an A record to a previously defined authoritative forward-mapping zone. If the zone is associated with one or more networks, the IP address must belong to one of the associated networks. For example, if the A record is in the corpxyz.com zone, which is associated with 10.1.0.0/16 network, then the IP addresses of the A record must belong to the 10.1.0.0/16 network. For information about associating zones and networks, see [Associating Networks with Zones](#).

The appliance also supports wildcard A records. For example, you can use a wildcard A record in the corpxyz.com domain to map queries for names such as www1.corpxyz.com, ftp.corpxyz.com, main.corpxyz.com, and so on to the IP address of a public-facing web server. Note that wildcard names only apply when the domain name being queried does not match any resource record.

NIOS allows superusers to add A records with a blank name. Limited-access users must have read/write permission to **Adding a blank A/AAAA record** to add A records with a blank name. You can assign global permission for specific admin groups and roles to allow limited-access users to add blank A records. For more information, see [Administrative Permissions for Adding Blank A or AAAA Records](#).



#### Note

If an A record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Name** field displays the record name in UTF-8 encoded format. For example, an A record with the domain name 工作站.test.com added through DDNS updates displays \229\183\165\228\189\156\231\171\153.test.com in the **Name** field.

#### Adding A Records

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add A Record**.
2. In the *Add A Record* wizard, do the following:
  - **Name:** If Grid Manager displays a zone name, enter the host name that you want to map to an IP address. The displayed zone name can either be the last selected zone or the zone from which you are adding the host record. If no zone name is displayed or if you want to specify a different zone, click **Select**



- Zone.** When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box and then enter the host name. The name you enter is prefixed to the DNS zone name that is displayed, and the complete name becomes the FQDN (fully qualified domain name) of the host. For example, if the zone name displayed is `corpxyz.com` and you enter `admin`, then the FQDN becomes `admin.corpxyz.com`. Ensure that the domain name you enter complies with the host name restriction policy defined for the zone. To create a wildcard A record, enter an asterisk (\*) in this field.
- **DNS View:** This field displays the DNS view to which the DNS zone belongs.
  - **Shared Record Group:** This field appears only when you are creating a shared record. Click **Select Shared Record Group**. If you have only one shared record group, the appliance displays the name of the shared record group here. If you have multiple shared record groups, select the shared record group in the *Shared Record Group Selector* dialog box. You can use filters or the **Go to** function to narrow down the list.
  - **Host Name Policy:** Displays the host name policy of the zone.
  - In the **IP Addresses** section, click the Add icon and do one of the following:
    - Select **Add Address** to enter the IPv4 address to which you want the domain name to map. or
    - Select **Next Available IPv4** to retrieve the next available IP address in a network.
    - If the A record is in zone that has associated networks, the *Network Selector* dialog box lists the associated networks. If the zone has no network associations, the *Network Selector* dialog box lists the available networks. When you select a network, Grid Manager retrieves the next available IP address in that network.
  - **Comment:** Optionally, enter additional information about the A record.
  - **Create associated PTR record:** Select this option to automatically generate a PTR record that maps the specified IP address to the host name. To create the PTR record, the reverse-mapping zone must be in the database.
  - **Disable:** Select this checkbox to disable the record. Clear the checkbox to enable it.
3. Click **Next** to define extensible attributes. For information, see [Managing Extensible Attributes](#).
  4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Modifying A Records

When you modify an A record, you can do the following:

- In the **General** tab, you can change the information you previously entered through the wizard.
- The **Discovered Data** tab displays discovered data, if any, for the record. For information, see [Viewing Discovered Data](#).

You can also enter or edit information in the **TTL**, **Extensible Attributes**, and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) below.

## Managing ALIAS Records

An ALIAS record is a virtual DNS record type created for a standard record type to alias the root domain (apex zone) to another name. You can use an ALIAS record to host a website at a domain name without the "www" (or other) prefix when using the cloud services, such as Amazon Web Services, Azure VMs, GitHub pages, Heroku, and so on. For example, you can use an ALIAS record to point your domain `foo.com` to a host name like `mail.foo.com`. When you perform a DNS lookup on an ALIAS record, the authoritative DNS server dynamically resolves an ALIAS record for the matching alias target. The response contains the aliased records, that can be A, AAAA, MX, NAPTR, PTR, SPF, SRV, or TXT values. You can alias the same domain with multiple target types.

You can synchronize ALIAS records from your AWS to NIOS using Amazon Route 53. For information about AWS deployments, refer to the *Installation Guide for vNIOS for AWS*.

Following are some guidelines to remember when you use ALIAS records:

- ALIAS records are not supported on DNS zones that are assigned to a Microsoft primary server. This means that you cannot assign a DNS zone containing an ALIAS record to a Microsoft primary server and you cannot add an ALIAS record to a DNS zone that is assigned to a Microsoft primary server.
- You cannot add an ALIAS record to a DNSSEC signed zone and you cannot sign a zone containing an ALIAS record.

- You cannot add an ALIAS record to a DNS zone if there is a secondary server that is using zone transfer as an update mechanism. Also, you cannot use zone transfer process to update zones containing ALIAS records in the Grid secondaries.
- You cannot update any ALIAS record using a DNS request. You can create and update ALIAS records only by using the NIOS UI or API.
- An ALIAS record can coexist with other types of records for the same owner name unless there is another configured record at the owner name whose type is an allowed target type for ALIAS. You cannot configure an ALIAS record whose target type is the same as the configured record type.
- You cannot use DDNS updates to add or delete an ALIAS record.

## Adding ALIAS Records

To add an ALIAS record, perform the following steps:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> ALIAS Record**.
2. In the *Add ALIAS Record* wizard, do the following:
  - **Name:** The ALIAS record name. The displayed zone name can either be the last selected zone or the zone from which you are adding the ALIAS record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. If you do not specify the record name, then it resumes the name of the Zone Apex.
  - **Record Type:** You can configure any record type - A, AAAA, MX, NAPTR, PTR, SPF, SRV, TXT.
  - **Target:** For for A, AAAA, and MX records, enter the host name as the target. For other record types, enter the domain name that is used to reply to any DNS request. You can also type the domain name for the resource. Examples:
    - CloudFront distribution domain name: d111111abcdef8.cloudfront.net
    - ELB load balancer DNS name: example-1.us-east-1.elb.amazonaws.com
    - S3 website endpoint: s3-website.us-east-2.amazonaws.com
    - Resource record set in this hosted zone: [www.example.com](http://www.example.com)
  - **Comment:** Enter additional information about the ALIAS record.
  - **Disable:** Select this checkbox to disable the record. Clear the checkbox to enable it.
3. Click **Next** to define extensible attributes. For information, see [Managing Extensible Attributes](#).
4. Save the configuration or click **Next** to schedule this task. Click **Now** in the **Schedule Change** panel to immediately execute this task or click **Later** and specify a date, time, and time zone. For information about how to schedule a task, see [Scheduling Tasks](#).
5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Modifying ALIAS Records

When you modify an ALIAS record, you can perform the following step:

- In the **General** tab, you can change the information you previously entered through the wizard.

You can also enter or edit information in the **TTL**, **Extensible Attributes**, and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) below.

## Managing NS Records

An NS record identifies an authoritative DNS server for a domain. Each authoritative DNS server must have an NS record. Grid Manager automatically creates NS records when you assign a Grid member as the primary server for a zone or when you assign an NS group to a zone. Grid Manager generates two NS records; an authoritative NS record for the current zone; and a delegation NS record for the parent zone for each name server available in the NS group. You cannot edit an automatically generated NS record. Note that when you delete a name server from an NS group, the NS record associated with the name server is deleted. For information about using NS Groups, see [Importing Zone Data](#). You can manually create NS records for other zones. NS records associated with one or more IP addresses are used for related A record and PTR record generation. You can configure an NS record for anycast IP addresses on the appliance. For more information about anycast, see [About Anycast Addressing for DNS](#).

## Adding NS Records

To add an NS record, perform the following steps:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add NS Record**.
2. In the *Add NS Record* wizard, complete the following fields:
  - **Zone:** The displayed zone name can either be the last selected zone or the zone from which you are adding the NS record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box.
  - **DNS View:** Displays the DNS view to which the selected zone belongs.
  - **Hostname Policy:** Displays the host name policy of the selected zone.
  - **Name Server:** Enter the host name that you want to configure as the name server for the zone. IDN is not supported in this field. You can use the punycode representation of an IDN in this field.
3. Click **Next** to enter IP addresses for the name server.
4. In the *Name Server Addresses* panel, click the Add icon and complete the following fields:
  - **Address:** Enter the IP address of the name server.
  - **Add PTR Record:** This field displays Yes by default, enabling the automatic generation of a PTR record for the IP address. You can select **No** to disable the generation of the PTR record.
5. Click **Next** to define extensible attributes or save the configuration and click **Restart** if it appears at the top of the screen.

## Modifying and Deleting NS Records

When you modify an NS record, you can change the following information:

- In the **General** tab, you can change the name server name.
- In the **Addresses** tab, you can do the following:
  - Delete an address by selecting it and clicking the Delete icon.
  - Add an address by clicking the Add icon, and then entering the IP address and completing the **Add PTR Record** field.

## Managing AAAA Records

An AAAA (quad A address) record maps a domain name to an IPv6 address. To define a specific name-to-address mapping, add an AAAA record to a previously defined authoritative forward-mapping zone. If the zone is associated with one or more networks, the IP address must belong to one of the associated networks. For example, if the AAAA record is in the corpxyz.com zone, which is associated with the 1111:0001/32 network, then the IP addresses of the A record must belong to that network. For information about associating zones and networks, see [Associating Networks with Zones](#).



### Note

If an AAAA record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Name** field displays the record name in UTF-8 encoded format. For example, an AAAA record with the domain name 工作站.test.com added through DDNS updates displays \229\183\165\228\189\156\231\171\153.test.com in the **Name** field.

NIOS allows superusers to add AAAA records with a blank name. Limited-access users must have read/write permission to **Adding a blank A/AAAA record** to add AAAA records with a blank name. You can assign global permission for specific admin groups and roles to allow limited-access users to add blank AAAA records. For more information, see [Administrative Permissions for Adding Blank A or AAAA Records](#).

## Adding AAAA Records

To create an AAAA record, perform the following steps:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add AAAA Record**.
2. In the *Add AAAA Record* wizard, complete the following:
  - **Name:** If Grid Manager displays a zone name, enter the host name that you want to map to an IP address. The displayed zone name can either be the last selected zone or the zone from which you are adding the AAAA record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter the host name. The name you enter is prefixed to the DNS zone name that is displayed, and the complete name becomes the FQDN (fully qualified domain name) of the host. For example, if the zone name displayed is corpxyz.com and you enter admin, then the FQDN becomes admin.corpxyz.com.
  - **DNS View:** Displays the DNS view to which the selected DNS zone belongs.
  - **Shared Record Group:** This field appears only when you are creating a shared record. Click **Select Shared Record Group**. If you have only one shared record group, the appliance displays the name of the shared record group here. If you have multiple shared record groups, select the shared record group in the *Shared Record Group Selector* dialog box. You can use filters or the **Go to** function to narrow down the list.
  - **Hostname Policy:** Displays the host name policy of the zone.
  - **IP Address:** Enter the IPv6 address to which you want the domain name to map. When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2006:0000:0000:0123:4567:89ab:0000:cdef can be shortened to 2006::123:4567:89ab:0:cdef. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The NIOS appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered.
  - **Comment:** Optionally, enter additional information about this record.
  - **Create associated PTR record:** Select this option to automatically generate a PTR record that maps the specified IP address to the host name. To create the PTR record, the reverse-mapping zone must be in the database.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes. For information, see [Managing Extensible Attributes](#).
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Modifying AAAA Records

When you modify an AAAA record, you can perform the following steps:

- In the **General** tab, you can change the information you previously entered through the wizard.
- In the **Discovered Data** tab, you can view discovered data, if any, for the record. For information, see [Viewing Discovered Data](#).

You can also enter or edit information in the **TTL**, **Extensible Attributes** and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) below.

## Managing PTR Records

In a reverse-mapping zone, a PTR (pointer) record maps an IP address to a host name. Before adding a PTR record to a reverse-mapping zone, you must first create the zone. You can also add PTR records to forward-mapping zones to support zeroconf (zero configuration networking), such as wide-area Bonjour. For information about the Bonjour protocol, refer to <http://www.apple.com/support/bonjour>. Though adding PTR records to forward-mapping zones supports some of the use cases in RFC 1101, it does not support the network name mapping use case described in the RFC. For more information, refer to <http://tools.ietf.org/html/rfc1101>.



#### Note

If a PTR record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Name** and **Domain Name** fields display the record name in UTF-8 encoded format. For example, a PTR record with the domain name 工作站.test.com added through DDNS updates displays \229\183\165\228\189\156\231\171\153.test.com in the **Name** and **Domain Name** fields.

## Adding PTR Records

To add a PTR record, perform the following steps:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add PTR Record**.
2. In the *Add PTR Record* wizard, do the following:
  - **Name or IP Address:** From the drop-down list, select **Name** or **IP Address**. When you select **Name**, click **Select Zone** to select a zone, and then enter a value for the PTR record. When you are adding a PTR record to a reverse-mapping zone, you can enter a value from 0 to 255 in the **Name** or **IP Address** field. Note that when you launch this wizard from the **IPAM** tab, you can only select a reverse-mapping zone. When you launch this from a reverse-mapping zone, the IP address field is populated with the prefix that corresponds to the selected zone. When you launch this from a forward-mapping zone, you can only specify the host name, not an IP address.
  - When you select **IP Address**, enter the IPv4 or IPv6 address that you want to map to the domain name.
  - **DNS View:** If you entered an IP address, you must select the DNS view of the PTR record. If you entered a name, this field displays the DNS view of the selected zone.
  - **Domain Name:** Enter the domain name to which you want the PTR record to point. For example, you can enter corpxyz.com.
  - **Comment:** Optionally, enter information about the PTR record.
  - **Disable:** Select this checkbox to disable the record. Clear the checkbox to enable it.
3. Save the configuration or click **Next** to define extensible attributes. For information, see [Managing Extensible Attributes](#).
4. Click **Restart** if it appears at the top of the screen.

To schedule this task, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, click **Later**, and then specify a date, time, and time zone.



#### Note

When you add a PTR record to a forward-mapping zone, a message may appear on the top of the wizard if a Grid member is configured to ignore DNS queries for PTR records in forward-mapping zones. Contact Infoblox Technical Support for more information about this message.

## Modifying PTR Records

To modify a PTR record, perform the following steps:

- In the **General** tab, you can change the information you previously entered through the wizard. Note that you cannot change an IPv4 address to an IPv6 address or move a PTR record from a forward-mapping zone to a reverse-mapping zone and vice versa. When you modify a PTR record that belongs to a forward-mapping zone, you can only modify the name since there is no IP address for such record.
- In the **Discovered Data** tab, you can view discovered data, if any, for the record. For information, see [Viewing Discovered Data](#).

You can also enter or edit information in the **TTL**, **Extensible Attributes**, and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) below.

## Managing MX Records

An MX (mail exchanger) record maps a domain name to a mail exchanger. A mail exchanger is a server that either delivers or forwards mail. You can specify one or more mail exchangers for a zone, as well as the preference for using each mail exchanger. A standard MX record applies to a particular domain or subdomain.

You can use a wildcard MX record to forward mail to one mail exchanger. For example, you can use a wildcard MX record in the corpxyz.com domain to forward mail for eng.corpxyz.com and sales.corpxyz.com to the same mail exchange, as long as the domain names do not have any matching resource record. Wildcards only apply when the domain name being queried does not match any record.



### Note

If an MX record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Mail Destination** and **Mail Exchanger** fields display the record name in UTF-8 encoded format. For example, an MX record with the domain name 工作站.test.com added through DDNS updates displays \229\183\165\228\189\156\231\171\153.test.com in the **Mail Destination** and **Mail Exchanger** fields.

## MX Records

The following MX records ... direct queries for one or more domains ... to the same mail exchanger:

An MX record for the mail exchanger that answers queries for just the corpxyz.com domain (and its corresponding A record):

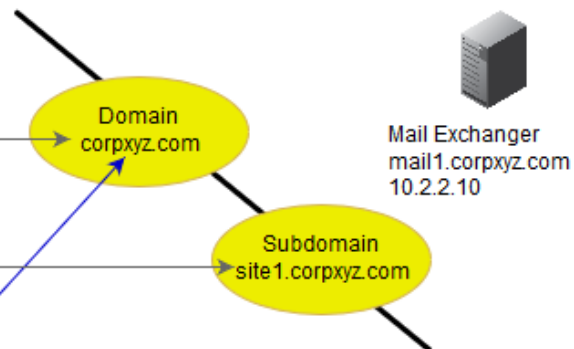
```
corpxyz.com IN MX 0 mail1.corpxyz.com  
mail1.corpxyz.com IN A 10.2.2.10
```

An MX record for just site1.corpxyz.com, a subdomain of corpxyz.com:

```
site1.corpxyz.com IN MX 0  
mail1.corpxyz.com
```

A wildcard MX record for the corpxyz.com domain:

```
*.corpxyz.com IN MX 0 mail1.corpxyz.com
```



### Note

You must also create an A record for the host defined as a mail exchanger in an MX record.

## Adding MX Records

To add an MX record from the Tasks Dashboard, see as described in [Add MX Record](#). You can also add MX records from the **Data Management** tab -> **DNS** tab by clicking **Add** -> **Record** -> **Add MX Record** from the Toolbar.

## Modifying and Deleting MX Records

When you modify an MX record, you can change the information you previously entered in the **General** tab. You can also enter or edit information in the **TTL**, **Extensible Attributes**, and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) below.



## Managing SRV Records

An SRV (service location) record directs queries to hosts that provide specific services. For example, if you have an FTP server, then you might create an SRV record that specifies the host which provides the service. You can specify more than one SRV record for a host. For more information about SRV records, see *RFC2052,ADNSRRforspecifyingthelocationofservices(DNSSRV)*.



### Note

If an SRV record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Name** and **SRV Target** fields display the domain name in UTF-8 encoded format. For example, an SRV record with the domain name 电脑.test.com added through DDNS updates displays \231\148\181\232\132\145.test.com in the **Name** and **SRV Target** fields.

## Adding SRV Records

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add SRV Record**.
2. In the *Add SRV Record* wizard, complete the following fields:
  - **Display input as:** Select the format in which you want the SRV record to be displayed. When you select RFC 2782 format, the appliance follows the `_service._protocol.name` format as defined in RFC 2782. When you select Free format, enter the entire name in the Domain field.
  - **Service:** Specify the service that the host provides. You can either select a service from the list or type in a service, if it is not on the list. For example, if you are creating a record for a host that provides FTP service, select `_ftp`. To distinguish the service name labels from the domain name, the service name is prefixed with an underscore. If the name of the service is defined at <http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml>, use that name. Otherwise, you can use a locally-defined name.
  - **Protocol:** Specify the protocol that the host uses. You can either select a protocol from the list or type in a protocol, if it is not on the list. For example, if it uses TCP, select `_tcp`. To distinguish the protocol name labels from the domain name, the protocol name is prefixed with an underscore.
  - **Domain:** If Grid Manager displays a zone name, enter the name here to define an SRV record for a host or subdomain. The displayed zone name can either be the last selected zone or the zone from which you are adding the SRV record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter the name to define the SRV record. The NIOS appliance prefixes the name you enter to the domain name of the selected zone. For example, if you want to create an SRV record for a web server whose host name is `www2.corpxyz.com` and you define the SRV record in the `corpxyz.com` zone, enter `www2` in this field. To define an SRV record for a domain whose name matches the selected zone, leave this field blank. The NIOS appliance automatically adds the domain name (the same as the zone name) to the SRV record. For example, if you want to create an SRV record for the `corpxyz.com` domain and you selected the `corpxyz.com` zone, leave this field blank.
  - **Preview:** After you have entered all the information, this field displays the FQDN, which is the concatenation of the Service, Protocol, and Domain fields.
  - **Shared Record Group:** This field appears only when you are creating a shared record. Click **Select Shared Record Group**. If you have only one shared record group, the appliance displays the name of the shared record group here. If you have multiple shared record groups, select the shared record group in the *Shared Record Group Selector* dialog box. You can use filters or the **Go to** function to narrow down the list.
  - **Priority:** Select or enter an integer from 0 to 65535. The priority determines the order in which a client attempts to contact the target host; the domain name host with the lowest number has the highest priority and is queried first. Target hosts with the same priority are attempted in the order defined in the **Weight** field.

- **Weight:** Select or enter an integer from 0 to 65535. The weight allows you to distribute the load between target hosts. The higher the number, the more that host handles the load (compared to other target hosts). Larger weights give a target host a proportionately higher probability of being selected.
  - **Port:** Specify the appropriate port number for the service running on the target host. You can use standard or nonstandard port numbers, depending on the requirements of your network. You can select a port number from the list or enter an integer from 0 to 65535.
  - **Target:** Enter the canonical domain name of the host (not an ALIAS); for example, www2.corpxyz.com. In addition, you need to define an A record mapping the canonical name of the host to its IP address.
  - **Comment:** Enter a descriptive comment for the record.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Save the configuration or click **Next** to define extensible attributes. For information, see [Managing Extensible Attributes](#).
  4. Click **Restart** if it appears at the top of the screen.

## Modifying and Deleting SRV Records

To modify an SRV record, perform the following steps:

- In the **General** tab, the **Display input as** field displays the format in which the SRV record was configured. For RFC 2782 format, the appliance matches the `{service.protocol.name}` format and displays the corresponding information in the Service and Protocol fields. If the appliance cannot match the service and protocol, it displays the entire name in the Domain field. For Free format, the entire name is displayed in the Domain field.



### Note

The appliance does not match the service and protocol names to exactly how they appear in the drop-down lists. It only checks whether the first two parts of the names start with an underscore. If the first two parts do not start with an underscore, the appliance assumes it is a free format. For example, `_abc._xyz.name` is considered as RFC 2782 format even though `_abc` is not in the **Service** drop-down list, and `_xyz` is not in the **Protocol** drop-down list. Grid Manager displays `_abc` in the **Service** field and `_xyz` in the **Protocol** field. On the other hand, `"abc.xyz.name"` is considered as a free format because the first two parts do not start with underscores, and Grid Manager displays this in its entirety in the **Domain** field.

You can also enter or edit information in the **TTL**, **Extensible Attributes**, and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) below.

## Managing TXT Records

A TXT (text record) record contains supplemental information for a host. For example, if you have a sales server that serves only North America, you can create a text record stating this fact. You can create more than one text record for a domain name.



### Note

If a TXT record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Name** field displays the domain name in UTF-8 encoded format. For example, a TXT record with the domain name `电脑.test.com` added through DDNS updates displays `\231\148\181\232\132\145.test.com` in the **Name** field.

## Using TXT Records for SPF

SPF (Sender Policy Framework) is an anti-forgery mechanism designed to identify spam e-mail. SPF fights e-mail address forgery and makes it easier to identify spam, worms, and viruses. Domain owners identify sending mail servers in DNS. SMTP receivers verify the envelope sender address against this information, and can distinguish legitimate mail from spam before any message data is transmitted.

SPF makes it easy for a domain to say, "I only send mail from these machines. If any other machine claims that I'm



sending mail from there, they're not valid." For example, when an AOL user sends mail to you, an email server that belongs to AOL connects to an email server that belongs to you. AOL uses SPF to publish the addresses of its email servers. When the message comes in, your email servers can tell if the server that sent the email belongs to AOL or not.

You can use TXT records to store SPF data that identifies what machines send mail from a domain. You can think of these specialized TXT records as *reverse MX records* that e-mail servers can use to verify if a machine is a legitimate sender of an e-mail.

### SPF Record Examples

```
corpxyz.com. IN TXT "v=spf1 mx -all"
corpxyz.net. IN TXT "v=spf1 a:mail.corpxyz.com -all" corpxyz.net. IN TXT
"v=spf1 include:corpxyz.com -all" corpxyz.net. IN TXT "v=spf1 mx -all
exp=getlost.corpxyz.com" corpxyz.com. IN TXT "v=spf1 include:corp200.com
-all"
```



#### Note

If an SPF record goes beyond the BIND limit of 255 characters, Infoblox recommends that you break up the record into two TXT records.

### Adding TXT Records

To add an TXT record from the Tasks Dashboard, see [Add TXT Record](#). You can also add TXT records from the **Data Management** tab -> **DNS** tab by clicking **Add** -> **Record** -> **Add TXT Record** from the Toolbar.

### Modifying and Deleting TXT Records

When you modify a TXT record, you can change the information you previously entered in the **General** tab. You can also enter or edit information in the **TTL**, **Extensible Attributes**, and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) below.

### Managing TLSA Records

A TLSA record is used to associate a TLS (Transport Layer Security) server certificate or a public key with a domain name. For example, you can define whether a certificate or a public key must be associated with a domain name when you define a TLSA resource record through Grid Manager. When you define your own TLSA record, you do not have to depend on an external Certificate Authority to issue a digitally signed TLS certificate for your domain name. When a client queries the domain name, TLSA records are matched to authenticate associated TLS certificates.

### Using TLSA Records for DANE

Infoblox supports DANE (DNS-based Authentication of Named Entities) protocol to secure information about domain names. DANE uses DNSSEC to sign certificates and keys that are used by the TLS and distributes secure information about the domain name using DNS. With DANE, you can make an authoritative binding between the domain name and a certificate or a public key, whichever is used by a host for the respective domain. You can define what kind of certificates or public keys must be associated with a domain name to prevent vulnerability attacks and to reduce or prevent the interaction of third-party Certification Authorities to issue PKIX certificates. For detailed information about the TLSA record format and certificate usage, refer to *RFC 6698 The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA*.

## Adding TLSA Records

In NIOS 8.5, you can add a TLSA record to a DNSSEC signed zone only. You cannot unsign a zone that contains a TLSA record. In NIOS 8.5.1 or later, you can add a TLSA record to a DNSSEC signed zone or an unsigned zone. To add a TLSA record:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> TLSA Record**.
2. In the *Add TLSA Record* wizard, complete the following fields:
  - **Display input as:** Select either **Strict format** (`_port._protocol.domain`) or **Free format**. Grid Manager selects **Strict format** by default. In this format, you can choose port and protocol values from the list. When you select **Free format**, you cannot specify these values.
  - **Port:** Select a value from the drop-down list to indicate the port on which the TLS-based service is active. The values in the drop-down list are:
    - 21 (FTP)
    - 22 (SSH)
    - 23 (Telnet)
    - 25 (SMTP)
    - 80 (HTTP)
    - 88 (Kerberos)
    - 389 (LDAP)
    - 443 (HTTPS)
    - 464 (KPASSWD)
    - 3268 (GC)
  - **Protocol:** Select a value from the drop-down list to indicate the protocol that is used for secure communication. The values in the drop-down list are:
    - `_msdcs`
    - `_sites`
    - `_tcp`
    - `_udp`When you select **Strict format**, **Port** and **Protocol** values are set to **443 (HTTPS)** and `_tcp`, by default. You can change these values. When you select **Free format**, you cannot edit the mentioned values.
  - **Name:** Enter a name for the TLSA resource record. You can specify a name only when you select **Free format**.
  - **Select Zone:** Click to select a zone. In NIOS 8.5, you must select only a signed zone to associate with a TLSA resource record. In NIOS 8.5.1 or later, you can select a signed zone or an unsigned zone. For more information, see [Signing a Zone](#). Click **Clear** to clear the **Name** that you have entered.
  - **FQDN:** This is displayed by default. You cannot modify the value. TLSA resource records are stored using the domain name that you select. When you select **Free format**, `name.domain` is displayed as the FQDN. Example: `abc.example.com`. When you select **Strict format**, `_port._protocol.domain` is displayed as the FQDN, where:
    - `_port` indicates the port on which the TLS-based service is active.
    - `_protocol` indicates the name of the transport protocol that you have selected.Consider an example where you are the owner of the domain `www.example.com` and you have set the **Port** to **443(HTTPS)** and **Protocol** to `tcp`, which indicates that the HTTP server is running TLS on port 443. To request TLSA record for `www.example.com`, you must use `__443._tcp.www.example.com`. Similarly, to request a TLSA resource record for an SMTP server running the STARTTLS protocol on port 25 at `mail.example.com`, you must use `_25._tcp.mail.example.com`.
  - **DNS View:** The DNS View associated with the selected DNS zone is displayed.
  - **Certificate Usage:** Select a value from the drop-down list to indicate how the certificate or the public key associated with the domain name is matched when the client queries for the domain name on the TLS server. The values in the drop-down list are: **PKIX-TA**, **PKIX-EE**, **DANE-TA**, and **DANE-EE**.
    - With **PKIX-TA** and **PKIX-EE**, you need additional Trust Anchors to validate peer certificate chains. These Trust Anchors must be mutually trusted by both the TLS server and the client. For more information, refer to *RFC 6698 The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA*.

- When you select **DANE-TA** and **DANE-EE**, the TLSA records that you define using Grid Manager are sufficient to verify the client's certificate chain and additional Trust Anchors are not required to authenticate the public key or certificate data. For more information, refer to *RFC 6698 The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA*.
- **Selector**: Select a value from the drop-down list to indicate whether you are associating an entire certificate or only the public key with the domain. When you select a value, it indicates which part of the TLS certificate presented by the server is matched with the associated data. The values in the drop-down list are **Full certificate** and **Subject Public Key Info**. NIOS builds a hexadecimal format for the entire certificate when you select **Full certificate**. If you select **Subject Public Key Info**, NIOS extracts the public key and builds a hexadecimal format for it.
- **Matched Type**: Select a value from the drop-down list to indicate how a TLS certificate or the public key of the domain received from the client must be matched with the certificate or the key that you have specified for the respective domain in the TLS server. You can select to match the entire content or only the hash of the selector. The values in the drop-down list are: **No hash**, **SHA 256 bit**, and **SHA 512 bit**. If you select **No hash**, the TLS server performs an exact match on the selected content. When you select either **SHA 256 bit** or **SHA 512 bit**, only the hash of the selected content is matched by the TLS server.
- **Certificate Data**: Enter the certificate data that must be matched for authentication. You can either paste the full certificate or the corresponding public key when the **Matched Type** is set to **No hash**. Based on the values that you select for the **Selector** and the **Matched Type**, the server builds a hexadecimal format for the TLSA record. If you set the **Matched Type** to **SHA 256 bit** or **SHA 512 bit**, you must specify only the hash of the full certificate or the public key.
- **Get From File**: Click this to upload the certificate or the public key to the server.  
Note the following:
  - When you select **Strict format**, you must provide either the certificate or public key or hash of any of them. The value must be based on the **Selector** and **Matched Type** field values.
  - When you select **Free format**, you must upload the certificate in DER format. The server builds an appropriate hexadecimal format for the TLSA record based on the **Selector** value.
- **Comment**: Optionally, enter a descriptive comment for the TLSA record.
- **Disable**: Clear the checkbox to enable the record. Select the checkbox to disable it.

You can also perform the following steps:

- Use **Global Search** to search for TLSA records. For information, see [Global Search](#).
- Use **Copy Records** to copy TLSA records between DNS zones. For information, see [Copying Zone Records](#).
- Define global permission for **All TLSA records** with read-only, read/write or deny access. You can also define object level permission for TLSA records. For information, see [Defining Global Permissions](#) and [Defining Object Permissions](#).
- Import and export records in CSV format. For information, see [Importing and Exporting Data using CSV Import](#).
- View audit log entries for the TLSA record. For information, see [Viewing the Audit Log](#).

## Modifying and Deleting TLSA Records

When you modify a TLSA record, you can change the information you previously entered in the **General** tab. You can also enter or edit information in the **TTL**, **Extensible Attributes**, and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) below.

## Managing CAA Records

A Certification Authority Authorization (CAA) DNS resource record enables domain owners to define the Certificate Authorities (CAs) that can issue certificates for a domain. When you define a CAA record, only the CAs listed in the records can issue certificates for the respective domain. With CAA, you can also define notification rules to manage requests for a certificate from a non-authorized CA. If you do not define a CAA resource record, any CA can issue a certificate for the domain. For detailed information about the CAA record format and certificate usage, refer to *RFC 6844 DNS Certification Authority Authorization (CAA) Resource Record*. You can add, edit, or delete a CAA resource record through Grid Manager or the Infoblox API. The following are a few examples of CAA resource records:

- example.com. CAA 0 issue "ssl.com; policy=ev"

- example.com. CAA 0 issuewild “;”
- example.com. CAA 0 iodef “mailto:certissues@example.com”
- example.com. CAA 0 iodef “certissues.example.com”



#### Note

When you enable threat protection on a member, you must configure either a pass rule or rate limiting rule for CAA DNS resource record types. This is specific to record types that use threat protection rule templates to allow incoming DNS queries for the respective CAA record. If you do not configure these rules, the threat protection service that is running on the member blocks the DNS queries of the CAA record.

## Adding CAA Records

To add a CAA record, perform the following steps:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> CAA Record**.
2. In the *Add CAA Record* wizard, complete the following fields:
  - **Name:** Enter a name for the CAA record. Click **Select Zone** to select a zone. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click **Clear** to clear the zone that you have entered.
  - **DNS View:** The DNS view associated with the selected DNS zone is displayed.
  - **Flag:** Select a checkbox to set the flag value. When the flag is set to **Bit 0 (Critical)**, it tells the CA that it must completely understand the property tag to proceed. A CA does not issue certificates for any domain when the flag is set to **Bit 0 (Critical)** and the property tag is not understood. NIOS considers the flag value as zero, if you do not select any checkbox.

Note that the flags are unsigned integers between 0 and 255. Infoblox represents these integers in the form of bits. When you select the checkbox for **Bit 0 (Critical)**, the flag value is set to binary 10000000, which is decimal 128. Example: CAA 128 xyz “Unknown”.

You can select only **Bit 0 (Critical)** as the flag value and the remaining checkboxes are reserved for future use. The appliance displays a warning message when you select a checkbox other than **Bit 0 (Critical)**.

Consider the following example with two CAA records:

- CAA 0 issue “ca.example.net; policy=ev”
- CAA 128 xyz “Unknown”

In the above example, the property tag **xyz** is flagged as unknown. The CA associated with example.net or any other issuer cannot issue a certificate unless the processing rules for the **xyz** property tag are clearly understood by the CA.

- **Type(Tag):** Indicates the type of CAA record. The supported CAA record types are:
  - **Issue:** Select this to explicitly authorize a single CA to issue a certificate for the domain and subdomains of the specified domain.
  - **Issuewild:** Select this to explicitly authorize a single CA to issue a wildcard certificate for the domain. It allows the domain holder or anyone acting under the authority of the domain holder to issue wildcard certificates for the domain.  
Note that **Issue wild** type takes precedence over **Issue**.
  - **Iodf:** Select this to specify an email address or URL of the web service to report invalid certificate requests or issued certificates that violate your CAA policy.  
Infoblox allows you to enter a new CAA record type other than those displayed in the drop-down list. The maximum length allowed is 255 characters.
- **Certificate Authority:** Indicates the CA that is authorized to issue a certificate for the domain. The maximum length for certificate authority is 8192 characters. You can also specify the email address or the URL to report CAA policy violation for the domain. This is valid for **Iodf** only. Infoblox recommends that you add either the **http://** or **https://** prefix to the domain name. You must explicitly add “mailto” when specifying the email address. For example, “mailto:admin@example.com”.
- **Comment:** Optionally, enter a descriptive comment for the CAA record.
- **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.

3. Save the configuration or click **Next** to define extensible attributes. For information, see [Managing Extensible Attributes](#).
4. Save the configuration or click **Next** to schedule this task. Click **Now** in the **Schedule Change** panel to immediately execute this task or click **Later** and specify a date, time, and time zone. For information about how to schedule a task, see [Scheduling Tasks](#).
5. Click **Save & Close** to complete the configuration.



#### Note

Infoblox does not support shared CAA records and does not provide Windows 2016 MS Server support for CAA records.

You can also perform the following steps:

- Use **Global Search** to search for CAA records. For information, see [Global Search](#).
- Use **Copy Records** to copy CAA records between DNS zones. For information, see [Copying Zone Records](#).
- Define global permission for **All CAA records** with read-only, read/write or deny access. You can also define object level permission for CAA records. For information, see [Defining Global Permissions](#) and [Defining Object Permissions](#).
- Import and export records in CSV format. For information, see [Importing and Exporting Data using CSV Import](#).
- View audit log entries for the CAA record. For information, see [Viewing the Audit Log](#).
- Use Smart Folders to organize threat protection profiles by name, comment or object type. For information, see [Smart Folders](#).
- You can view the status of the import process and a summary report in the Data Import Wizard Log. For large data sets, this option is an efficient approach. To download the Data Import Wizard, visit <https://data-import-wizard.infoblox.com/#/overviewDashboard>.

## Modifying and Deleting CAA Records

When you modify a CAA record, you can change the information you previously entered in the **General** tab. You can also enter or edit information in the **TTL**, **Extensible Attributes**, and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) below.

## Managing CNAME Records

A CNAME record maps an ALIAS to a canonical name. You can use CNAME records in both forward- and IPv4 reverse-mapping zones to serve two different purposes. (At this time, you cannot use CNAME records with IPv6 reverse-mapping zones.)



#### Note

If a CNAME record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **ALIAS** and **Canonical Name** fields display the domain name in UTF-8 encoded format. For example, a CNAME record with the domain name 电脑.test.com added through DDNS updates displays \231\148\181\232\132\145.test.com in the **Canonical Name** and **ALIAS** fields.

## CNAME Records in Forward-Mapping Zones

In a forward-mapping zone, a CNAME record maps an ALIAS to a canonical (or official) name. CNAME records are often more convenient to use than canonical names because they can be shorter or more descriptive. For example, you can add a CNAME record that maps the ALIAS *qa.engr* to the canonical name *qa.engr.corpxyz.com*.



#### Note

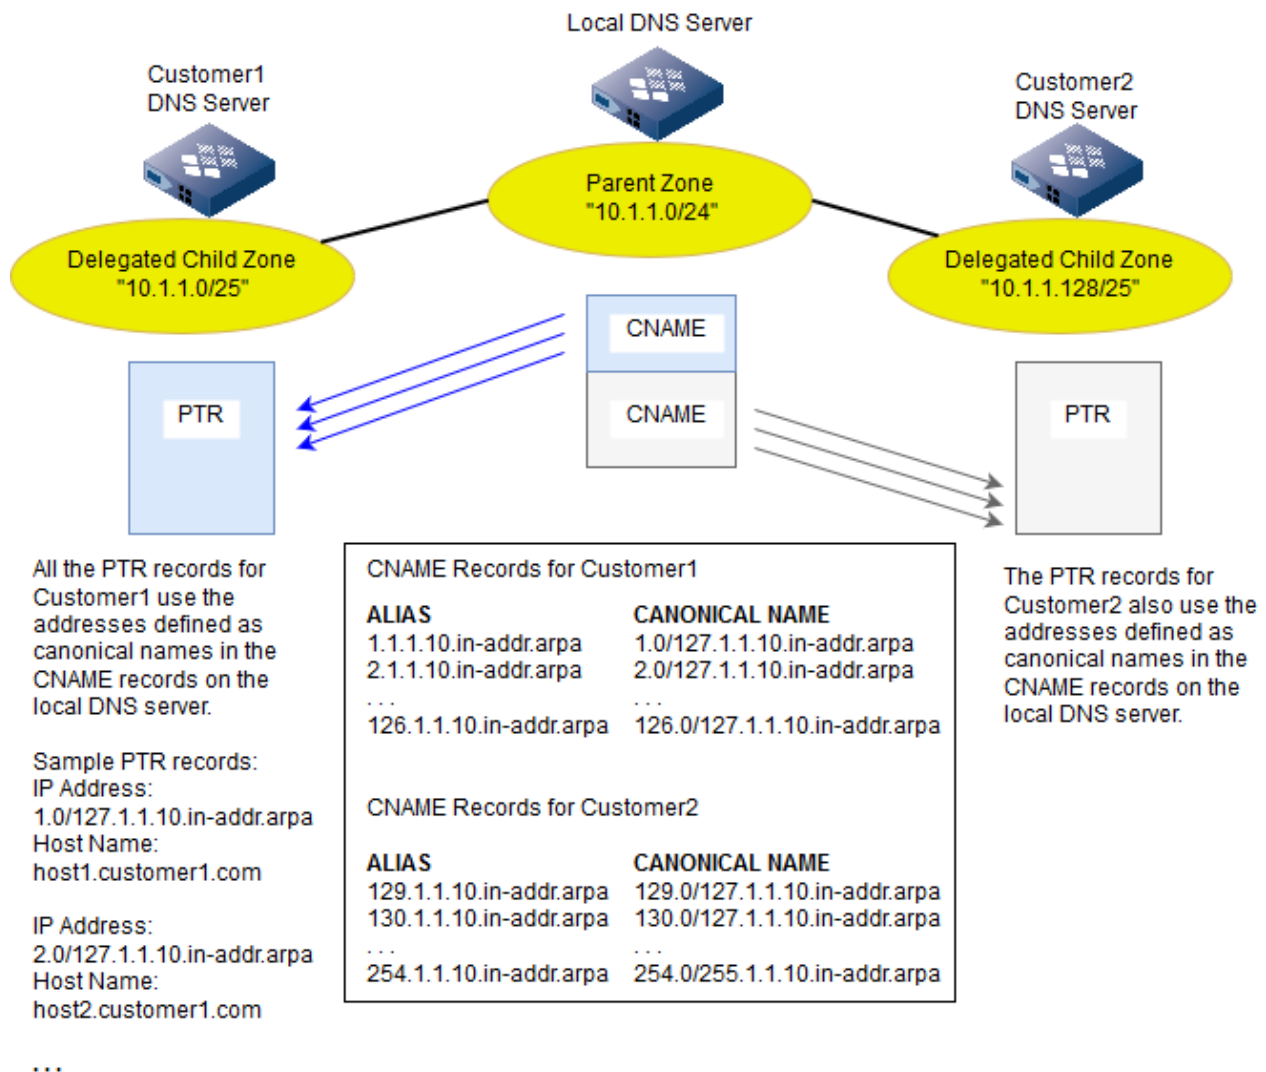
A CNAME record does not have to be in the same zone as the canonical name to which it maps. In addition, a CNAME record cannot have the same name as any other record in that zone.

To add a CNAME record to a forward-mapping zone from the Tasks Dashboard, see [Add CNAME Record](#). You can also add CNAME records from the **Data Management** tab -> **DNS** tab by clicking **Add** -> **Record** -> **Add CNAME Record** from the Toolbar.

### CNAME Records in IPv4 Reverse-Mapping Zones

You can add CNAME records to an IPv4 reverse-mapping zone to create ALIASes to addresses maintained by a different name server when the reverse-mapping zone on the server is a delegated child zone with fewer than 256 addresses. This technique allows you to delegate responsibility for a reverse-mapping zone with an address space of fewer than 256 addresses to another authoritative name server. See the following figure and *RFC 2317, Classless IN-ADDR.ARPA delegation*.

*CNAME Records in a Reverse-Mapping Zone*



You add CNAME records in the parent zone on your name server. The ALIASes defined in those CNAME records point to the addresses in PTR records in the child zone delegated to the other server.



When you define a reverse-mapping zone that has a netmask from /25 (255.255.255.128) to /31 (255.255.255.254), you must include an RFC 2317 prefix. This prefix can be anything, from the address range (examples: 0-127, 0/127) to descriptions (examples: first-network, customer1). On a NIOS appliance, creating such a reverse-mapping zone automatically generates all the necessary CNAME records. However, if you need to add them manually to a parent zone that has a child zone with fewer than 255 addresses.

### Adding CNAME Records

To add a CNAME record to a forward-mapping or reverse-mapping zone from the Tasks Dashboard, see [Add CNAME Record](#). You can also add CNAME records from the **Data Management** tab -> **DNS** tab by clicking **Add** -> **Record** -> **Add CNAME Record** from the Toolbar.

### Modifying and Deleting CNAME Records

When you modify a CNAME record, you can change the information you previously entered in the **General** tab. You can also enter or edit information in the **TTL**, **Extensible Attributes**, and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) below.

### Managing DNAME Records

A DNAME record maps all the names in one domain to those in another domain, essentially substituting one domain name suffix with the other (see RFC 2672, *Non-Terminal DNS Name Redirection*). For example, adding a DNAME record to the corpxyz.com domain mapping "corpxyz.com" to "corp200.com" maps *name-x.corpxyz.com* to *name-x.corp200.com*:

Domain Name		Target Domain Name
server1.corpxyz.com	—>	server1.corp200.com
server2.corpxyz.com	—>	server2.corp200.com
server3.corpxyz.com	—>	server3.corp200.com
...corpxyz.com	—>	...corp200.com



#### Note

If a DNAME record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **ALIAS** and **Target** fields display the domain name in UTF-8 encoded format. For example, a DNAME record with the domain name 电脑.test.com added through DDNS updates displays \231\148\181\232\132\145.test.com in the **ALIAS** and **Target** fields.

When a request arrives for a domain name to which a DNAME record applies, the NIOS appliance responds with a CNAME record that it dynamically creates based on the DNAME definition. For example, if there is a DNAME record:

```
corpxyz.com.
```

```
DNAME corp200.com.
```

and a request arrives for server1.corpxyz.com, the NIOS appliance responds with the following CNAME record:

server1.corpxyz.com.

CNAME

server1.corp200.com.

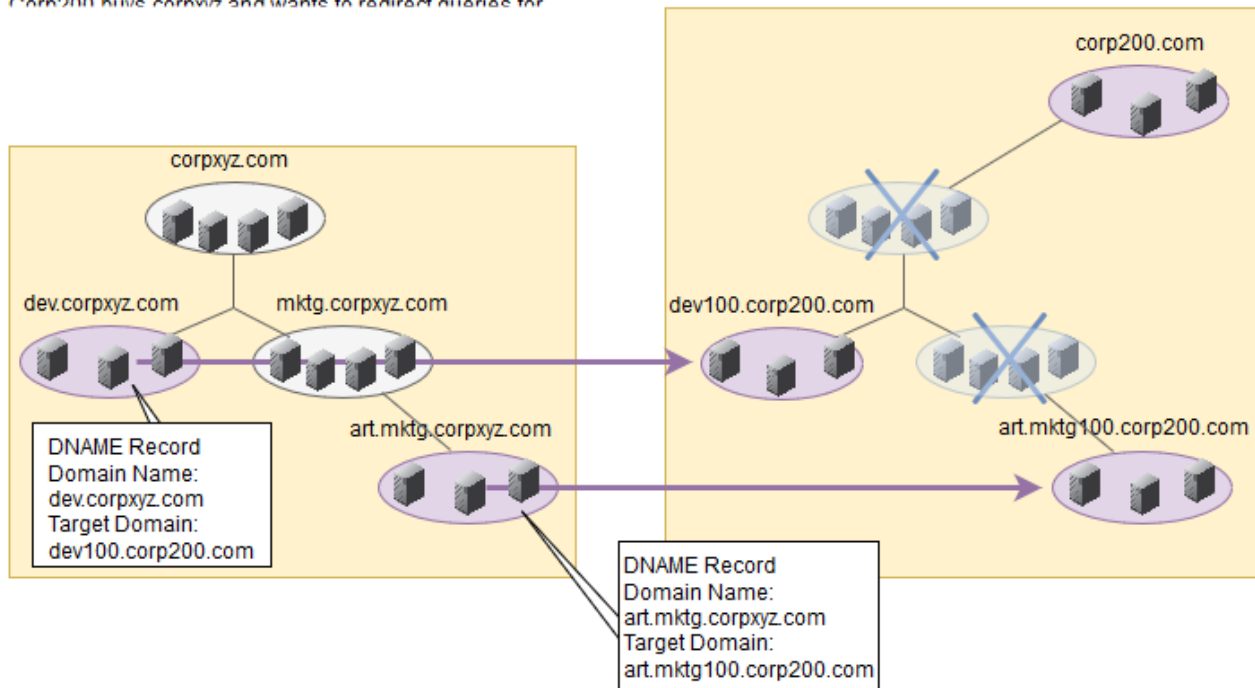
If responding to a name server running BIND 9.0.0 or later, the NIOS appliance also includes the DNAME record in its response, so that name server can also create its own CNAME records based on the cached DNAME definition.

The following are two common scenarios for using DNAME records:

- One company buys another and wants people using both the old and new name spaces to reach the same hosts.
  - A virtual Web hosting operation offers different "vanity" domain names that point to the same server or servers.
- There are some restrictions that apply to the use of DNAME records:
- You cannot have a CNAME record and a DNAME record for the same subdomain.
  - You cannot use a DNAME record for a domain or subdomain that contains any subdomains. You can only map the lowest level subdomains (those that do not have any subdomains below them). For an example of using DNAME records in a multi-tiered domain structure, see the following figure.

#### Adding DNAME Records for the Lowest Level Subdomains

Corp200 buys corpxyz and wants to redirect queries for

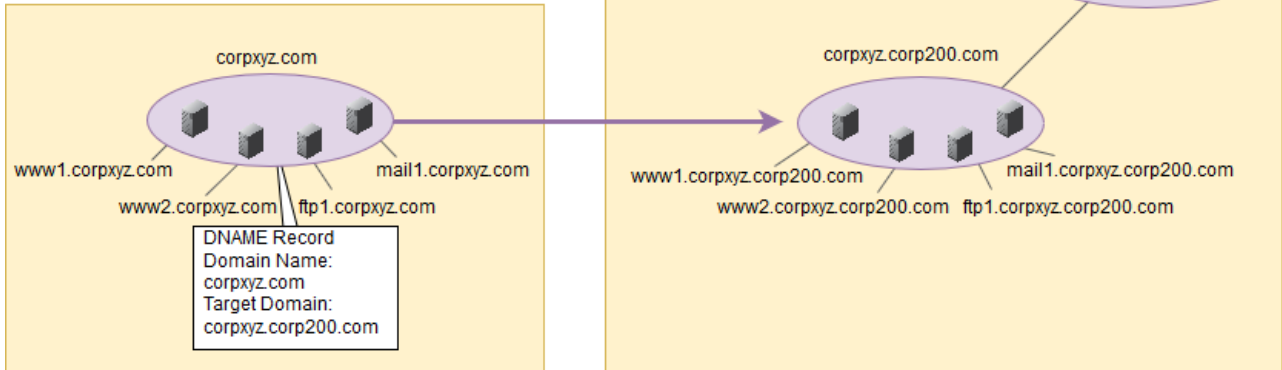


In the case of a domain structure consisting of a single domain (no subdomains), adding a DNAME record redirects queries for every name in the domain to the target domain, as shown in the following figure.

#### Adding a DNAME Record for a Single Domain



Corp200 buys corpxyz and wants to redirect all queries for corpxyz.com to corp200.com. To accomplish this, you add a single DNAME record to corpxyz.com.



When using a DNAME record, you must copy the resource records for the source domain to the zone containing the target domain, so that the DNS server providing service for the target domain can respond to the redirected queries.

Copy from corpxyz.com	to corpxyz.corp200.com
www1 IN A 10.1.1.10	www1 IN A 10.1.1.10
www2 IN A 10.1.1.11	www2 IN A 10.1.1.11
ftp1 IN A 10.1.1.20	ftp1 IN A 10.1.1.20
mail1 IN A 10.1.1.30	mail1 IN A 10.1.1.30

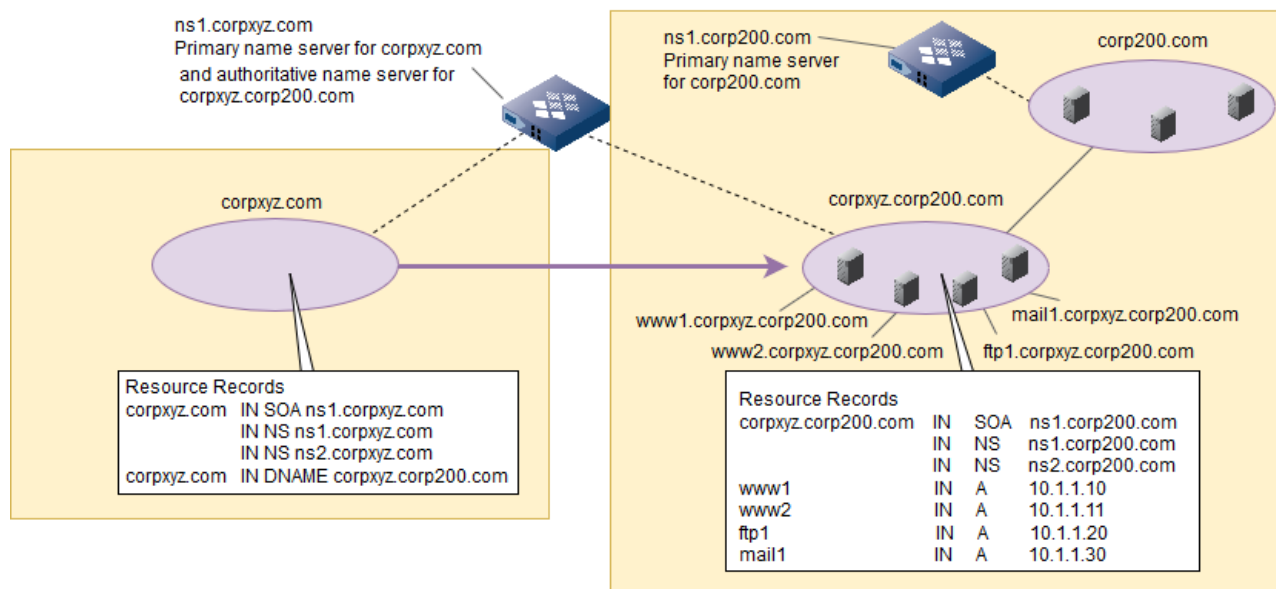
After copying these records to the zone containing the corpxyz.corp200.com domain, delete them from the zone containing the corpxyz.com domain.

If DNS service for the source and target domain names is on different name servers, you can import the zone data from the NIOS appliance hosting the source domain to the appliance hosting the target domain. For information about this procedure, see [Importing Zone Data](#).

If DNS service for the source and target domain names is on the same name server and the parent for the target domain is on a different server, you can delegate DNS services for the target domain name to the name server that provided—and continues to provide—DNS service for the source domain name (see the figure below). By doing this, you can continue to maintain resource records on the same server, potentially simplifying the continuation of DNS administration.

#### *Making the Target Zone a Delegated Zone*

On the primary name server for corp200.com (ns1.corp200.com), specify "corpxyz.corp200.com" as a delegated zone and specify "ns1.corpxyz.com" as the name server for that zone.



**Note:** This is a conceptual representation of domain name mapping and depicts the resulting hierarchical relationship of corp200.com as the parent zone for corpxyz.corp200.com. The hosts are not physically relocated.



#### Note

This is a conceptual representation of domain name mapping and depicts the resulting hierarchical relationship of corp200.com as the parent zone for corpxyz.corp200.com. The hosts are not physically relocated.

The following tasks walk you through configuring the two appliances in the *Making the Target Zone a Delegated Zone* figure to redirect queries for corpxyz.com to corpxyz.corp200.com using a DNAME record:

On the ns1.corpxyz.com name server, perform the following steps:

1. Create a new forward-mapping zone called corpxyz.corp200.com. See [Creating an Authoritative Forward-Mapping Zone](#).
2. Copy all the resource records for the domain or subdomain to which the DNAME record is going to apply from corpxyz.com to corpxyz.corp200.com.  
Because you can only specify the records by type, not individually, you might have to copy some records that you do not want and then delete them from the [corpxyz.corp200.com](#) zone.
3. In the corpxyz.com zone, delete all the resource records for the domain or subdomain to which the DNAME record is going to apply.
4. Add a DNAME record to the corpxyz.com zone specifying "corpxyz.com" as the domain and "corpxyz.corp200.com" as the target domain. Adding a DNAME record is explained in the next section.
5. On the ns1.corp200.com name server, add corpxyz.corp200.com as a delegated zone and specify ns1.corpxyz.com as the name server for it. See [Configuring a Delegation](#).

#### DNAME Records for Forward-Mapping Zones

To add a DNAME record to a forward-mapping zone:

- From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add DNAME Record**.

In the *Add DNAME Record* wizard, complete the following fields:



**Note**

If you specify a subdomain in the Domain Name field when configuring a DNAME record and the subdomain is also a subzone, the DNAME record appears in the list view for the subzone, not in the list view for the parent zone selected in the process of adding the record.

- **ALIAS:** If Grid Manager displays a zone name, enter the name of a subdomain here. If you are adding a DNAME record for the entire zone, leave this field empty. This field is for adding a DNAME record for a subdomain within the selected zone. The displayed zone name can either be the last selected zone or the zone from which you are adding the CNAME record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter the name of a subdomain.
- **Target:** Enter the domain name to which you want to map all the domain names specified in the ALIAS field.
- **Comment:** Enter identifying text for this record, such as a meaningful note or reminder.
- **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
- Save the configuration or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).
- Click **Restart** if it appears at the top of the screen.

### DNAME Records for Reverse-Mapping Zones

You can use DNAME records to redirect reverse lookups from one reverse-mapping zone to another. You can use DNAME records for reverse-mapping zones to simplify the management of subzones for classless address spaces larger than a class C subnet (a subnet with a 24-bit netmask).

RFC 2672, *Non-Terminal DNS Name Redirection*, includes an example showing the delegation of a subzone for an address space with a 22-bit netmask inside a zone for a larger space with a 16-bit netmask:

\$ORIGIN 0.192.in-addr.arpa.

8/22	NS	ns.slash-22-holder.example.
8	DNAME	8.8/22
9	DNAME	9.8/22
10	DNAME	10.8/22
11	DNAME	11.8/22

The reverse-mapping zone 0.192.in-addr.arpa. applies to the address space 192.0.0.0/16. Within this zone is a subzone and subdomain with the abbreviated name *8/22*. (Its full name is *8/22.0.192.in-addr.arpa*.) This subdomain contains its own subdomains corresponding to the 1024 addresses in the 192.0.8.0/22 subnet:

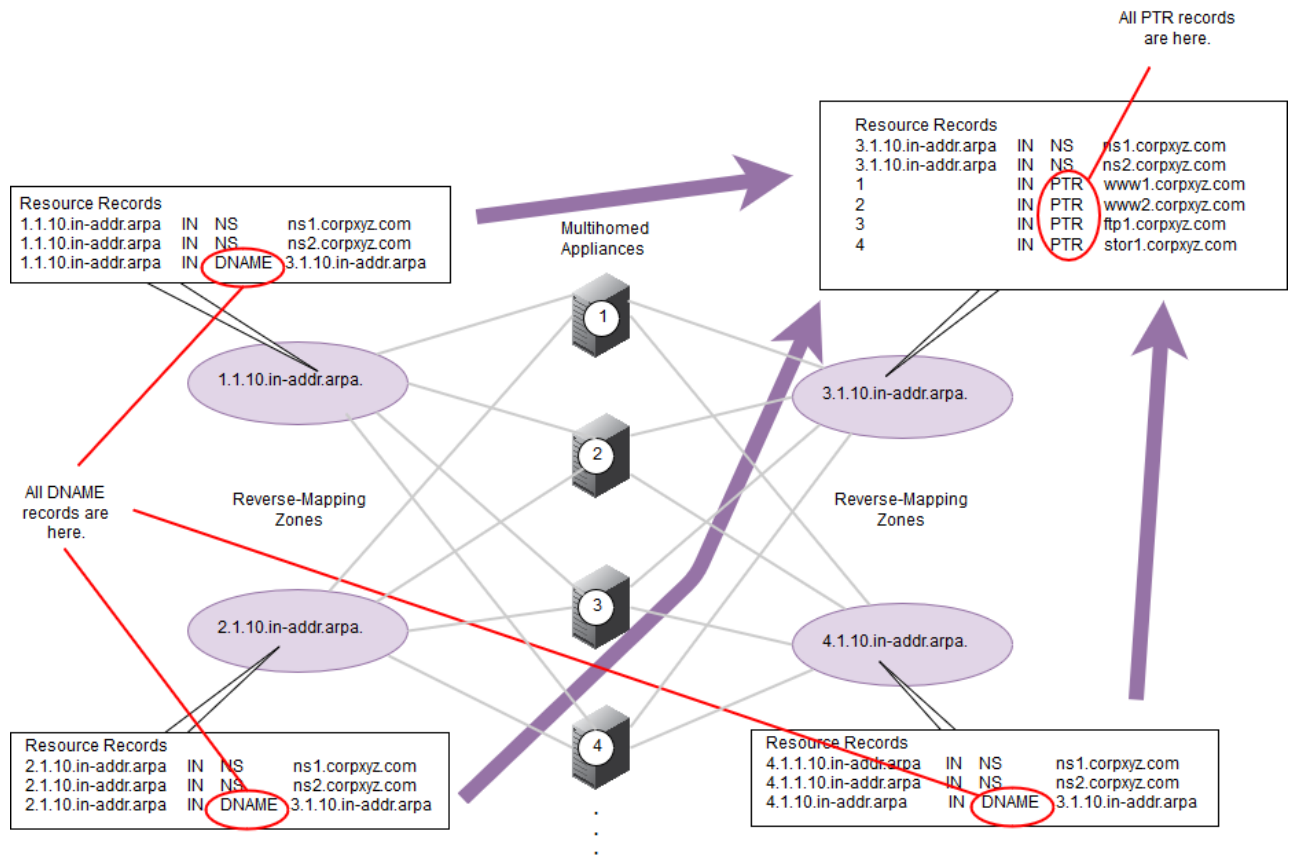
- Subdomain *8/22* (*8/22.0.192.in-addr.arpa*)
  - Subdomain *8.8/22* for addresses 192.0.8.0 – 192.0.8.255 (or 192.0.8.0/24)
  - Subdomain *9.8/22* for addresses 192.0.9.0 – 192.0.9.255 (or 192.0.9.0/24)
  - Subdomain *10.8/22* for addresses 192.0.10.0 – 192.0.10.255 (or 192.0.10.0/24)
  - Subdomain *11.8/22* for addresses 192.0.11.0 – 192.0.11.255 (or 192.0.11.0/24)

The NS record delegates authority for the reverse-mapping subzone 8/22 to the DNS server ns.slash-22-holder.example. Finally, the DNAME records provide ALIASes mapping domain names that correspond to the 192.0.8.0/24, 192.0.9.0/24, 192.0.10.0/24, and 192.0.11.0/24 subnets to the respective subdomains 8.8/22, 9.8/22, 10.8/22, and 11.8/22 in the 8/22.0.192.in-addr.arpa subzone.

**Note**  
 NIOS appliances support DNAME records in reverse-mapping zones that map addresses to target zones with a classless address space larger than a class C subnet. However, NIOS appliances do not support such target zones.

You might also use DNAME records if you have a number of multihomed appliances whose IP addresses must be mapped to a single set of domain names. An example of this is shown in the below figure.

*DNAME Records to Simplify DNS for Multihomed Appliances*



To add a DNAME record to a reverse-mapping zone, perform the following steps:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add DNAME Record**.
2. In the *Add DNAME Record* wizard, complete the following fields:
  - If you specify a subdomain in the Domain Name field when configuring a DNAME record and the subdomain is also a subzone, the DNAME record appears in the list view for the subzone, not in the list view for the parent zone selected in the process of adding the record.
    - **ALIAS:** If Grid Manager displays a zone name, enter the name of a subdomain here. If you are adding a DNAME record for the entire zone, leave this field empty. This field is for adding a DNAME record for a subdomain within the selected zone. The displayed zone name can either be the last selected zone or the zone from which you are adding the CNAME record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter the name of a subdomain.

- **Target:** Type the name of the reverse-mapping zone to which you want to map all the addresses specified in the Domain Name field.
  - **Comments:** Enter identifying text for this record, such as a meaningful note or reminder.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Save the configuration or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).
  4. Click **Restart** if it appears at the top of the screen.

### Modifying and Deleting DNAME Records

When you modify a CNAME record, you can change the information you previously entered in the **General** tab. You can also enter or edit information in the **TTL**, **Extensible Attributes** and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) below.

### Managing NAPTR Records

A NAPTR (Name Authority Pointer) record specifies a rule that uses a substitution expression to rewrite a string into a domain name or URI (Uniform Resource Identifier). A URI is either a URL (Uniform Resource Locator) or URN (Uniform Resource Name) that identifies a resource on the Internet.


NAPTR records are usually used to map E.164 numbers to URIs or IP addresses. An E.164 number is a telephone number, 1-555-123-4567 for example, in a format that begins with a country code, followed by a national destination code and a subscriber number. (E.164 is an international telephone numbering system recommended by the International Telecommunication Union.) Thus, NAPTR records allow us to use telephone numbers to reach devices, such as fax machines and VoIP phones, on the Internet.

To map an E.164 to a URI, the E.164 number must first be transformed into a domain name. ENUM (E.164 Number Mapping) specifies a method for converting E.164 numbers to domain names. For example, using the method specified by ENUM, the telephone number 1-555-123-4567 becomes the domain name 7.6.5.4.3.2.1.5.5.5.1.e164.arpa. For details about ENUM, refer to *RFC 3761, The E.164 to Uniform Resource Identifiers (URI) Dynamic Delegation Discovery System (DDDS) Application (ENUM)*.

After the E.164 number is converted to a domain name, a DNS client can then perform a DNS lookup for the NAPTR records of the domain name. The following example illustrates how a DNS client processes NAPTR records.

In this example, the telephone number 1-555-123-4567 is converted to the domain name 7.6.5.4.3.2.1.5.5.5.1.e164.arpa. The DNS client then sends a query to the Infoblox DNS server for the NAPTR records associated with 7.6.5.4.3.2.1.5.5.5.1.e164.arpa. The Infoblox DNS server returns the following NAPTR record:

```
$ORIGIN 7.6.5.4.3.2.1.5.5.5.1.e164.arpa
IN NAPTR 10 100 "U" "sip + E2U" "!^.*$!sip:jdoe@corpxyz.com!" .
```



The DNS client then examines the fields in the NAPTR record as follows:

- If a DNS client receives multiple NAPTR records for a domain name, the value in the Order field determines which record is processed first. It processes the record with the lowest value first.
- The DNS client uses the Preference value when the Order values are the same. Similar to the Preference field in MX records, this value indicates which NAPTR record the DNS client should process first when the records have the same Order value. It processes the record with the lowest value first. In the example, the DNS client ignores the Order and Preference values because it received only one NAPTR record.
- The Flag field indicates whether the current lookup is terminal; that is, the current NAPTR record is the last NAPTR record for the lookup. It also provides information about the next step in the lookup process. The flags that are currently used are:

**U:** Indicates that the output maps to a URI (Uniform Record Identifier).

**S:** Indicates that the output is a domain name that has at least one SRV record. The DNS client must then send a query for the SRV record of the resulting domain name.

**A:** Indicates that the output is a domain name that has at least one A or AAAA record. The DNS client must then send a query for the A or AAAA record of the resulting domain name.

**P:** Indicates that the protocol specified in the Service field defines the next step or phase.

- If the Flag field is blank, this indicates that the client must use the resulting domain name to look up other NAPTR records.
- The Service field specifies the service and protocol that are used to communicate with the host at the domain name. In the example, the service field specifies that SIP (Session Initiation Protocol) is used to contact the telephone service.
- The regular expression specifies the substitution expression that is applied to the original string of the client. In the example, the regular expression `!^.*$!sip:jd@corpxyz.com!` specifies that the domain name `7.6.5.4.3.2.1.5.5.1.e164.arpa` is replaced with `sip:jd@corpxyz.com`. The regular expression in a NAPTR record is always applied to the original string of the client. It must not be applied to a domain name that resulted from a previous NAPTR rewrite.
- The Replacement field specifies the FQDN for the next lookup, if it was not specified in the regular expression.



#### Note

If a NAPTR record with the domain name in its native characters is added to the Infoblox Grid through DDNS updates, the **Domain** and **Replacement** fields display the domain name in UTF-8 encoded format. For example, a NAPTR record with the domain name `电脑.test.com` added through DDNS updates displays `\231\148\181\232\132\145.test.com` in the **Domain** and **Replacement** fields.

## Adding NAPTR Records

To add a NAPTR record, perform the following steps:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Add NAPTR Record**.
2. In the *Add NAPTR Record* wizard, complete the following fields:
  - **Domain:** If Grid Manager displays a zone name, enter the domain name to which this resource record refers. The displayed zone name can either be the last selected zone or the zone from which you are adding the NAPTR record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter a domain name for the record. The name you enter is prefixed to the DNS zone name that is displayed, and the complete name becomes the FQDN (fully qualified domain name) of the record. For example, if the zone name displayed is `corpxyz.com` and you enter `admin`, then the FQDN becomes `admin.corpxyz.com`. This field is not displayed when you configure a NAPTR record for a DTC server.
  - **DNS View:** Displays the DNS view of the selected zone.
  - **Service:** Specifies the service and protocol used to reach the domain name that results from applying the regular expression or replacement. You can enter a service or select a service from the list.
  - **Flags:** The flag indicates whether the resulting domain name is the endpoint URI or if it points to another record. Select one of the following:
    - U:** Indicates that the output maps to a URI.
    - S:** Indicates that the resulting domain name has at least one SRV record.
    - A:** Indicates that the resulting domain name has at least one A or AAAA record.
    - P:** Indicates that this record contains information specific to another application. Leave this blank to indicate that the DNS client must use the resulting domain name to look up other NAPTR records. You can use the NAPTR records as a series of rules that are used to construct a URI or domain name.
  - **Order:** Select an Integer from 10 to 100, or enter a value from 0 to 65535. This value indicates the order in which the NAPTR records must be processed. The record with the lowest value is processed first.

- **Preference:** Select an Integer from 10 to 100, or enter a value from 0 to 65535. Similar to the Preference field in MX records, this value indicates which NAPTR record should be processed first when the records have the same Order value. The record with the lowest value is processed first.
  - **REGEX:** The regular expression that is used to rewrite the original string from the client into a domain name. RFC 2915 specifies the syntax of the regular expression. Note that the appliance validates the regular expression syntax between the first and second delimiter against the Python re module, which is not 100% compatible with POSIX Extended Regular Expression as specified in the RFC. For information about the Python re module, refer to <http://docs.python.org/release/2.5.1/lib/module-re.html>.
  - **Replacement:** This specifies the domain name for the next lookup. The default is a dot (.), which indicates that the regular expression in the **REGEX** field provides the replacement value. Alternatively, you can enter the replacement value in FQDN format.
  - **Comment:** Optionally, enter a descriptive comment for this record.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#). This is not applicable when you configure a NAPTR record for a DTC server.
  4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Managing LBDN Records

When your Grid has a DNS Traffic Control license, you can add LBDN (Load Balanced Domain Name) records to authoritative or delegated zones. You can add an LBDN even if the zone is DNSSEC signed but some restrictions apply. To add an LBDN record when in the DNS records list view:

- On the **Zones** or **Members/Servers** tab, click the arrow next to the Add icon and select **Record -> DTC LBDN**. For more information, see [Configuring DNS Traffic Control LBDNs](#).

You can also add an LBDN when in the Traffic Control tab. For more information, see the previously referenced section.

## Managing Unknown Records

An unknown record is a resource record for which the record type is unknown to NIOS. You can create an unknown record and assign a type to it. NIOS converts the unknown record to the type you assign. You can assign a type listed at <https://www.iana.org/assignments/dns-parameters/dns-parameters.xhtml>. For information about types that you cannot assign, see the *Guidelines for Creating Unknown Records* section.

## Adding Unknown Records

To add a record of Unknown type, perform the following steps:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Add -> Record -> Unknown Record**.
2. In the *Add Unknown Record* wizard, complete the following fields:
  - **Domain Name:** Click **Select Zone** to select a zone. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click **Clear** to clear the zone that you have entered.
  - **DNS View:** The DNS view associated with the selected DNS zone is displayed.
  - **Type:** Enter the type that the unknown record belongs to. You can either enter the type in the textual mnemonic format or in the "TYPE $nnn$ " format where "nnn" indicates the numeric type value. For example, for a record of type RP, you can either enter "RP" or "TYPE17".
3. Click the + icon to specify the details for the record you are creating:
  - **Field Type:** Select the field type that the record data must assume. Field types can be of the following:
    - **Base64-encoded Data:** BASE64 encoded binary data
    - **Hexadecimal Sequence:** Hexadecimal encoded binary data
    - **8-bit unsigned integer:** Unsigned 8-bit integer
    - **16-bit unsigned integer:** Unsigned 16-bit integer
    - **32-bit unsigned integer:** Unsigned 32-bit integer
    - **IPv4 Address:** IPv4 address in numerical form. For example, 192.0.1.1
    - **IPv6 Address:** IPv6 address in numerical form. For example, 2001:db8::abcd
    - **ASCII String:** ASCII text



- **Domain Name:** Domain name
  - **Presentation:** Standard textual form of record data, as shown in a standard master zone file. This type is specifically intended to be used for standard types of records that cannot easily be represented as a sequence of fields of the other types. Such record types include LOC and APL. If you choose this field type, it must be the only field to represent the record
  - **Value:** Value of the field data. Before entering a value, see the *Guidelines for Creating Unknown Records* section.
  - **Length:** Format in which to specify the length of the field value. The length can only be **None** for fields of 8-bit unsigned integer, 16-bit unsigned integer, 32-bit unsigned integer, IPv4 Address, IPv6 Address, Domain Name, and Presentation types. For fields of type Base64-encoded Data, ASCII String, and Hexadecimal Sequence, the value of the **Length** field can be either **None** or **8 bits** or **16 bits** depending on the requirement of the corresponding record type.  
Irrespective of the field type you select, there is an implementation-specific limitation on the length of the record data. Specifically, the data is internally converted to a textual form that appears in a standard DNS master file, and it is rejected if the converted text exceeds 8192 bytes. Although unlikely, some extremely large data can be rejected because of this limitation.
4. Click **Add**. The record details are added to the table below.
  5. In the **Comment** field, optionally enter a descriptive comment for the record.
  6. Clear the **Disable** checkbox to enable the record. Select the checkbox to disable it.
  7. Save the configuration or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).
  8. Click **Save & Close**.

## Guidelines for Creating Unknown Records

Make note of the following guidelines when you create an unknown record:

- You cannot enter a record type that already exists in NIOS. For example, A, AAAA, ANY, CAA, CNAME, DHCID, DNAME, DNSKEY, DS, MX, NAPTR, NS, NSEC, NSEC3, NSEC3PARAM, PTR, RRSIG, SOA, SRV, TLSA, TXT.
- If the record contains an ASCII String field type and you include double quotes, you must escape it with a backslash. For example, to obtain a value of "a"b", specify the string as \"a\"b\".
- Ensure that you use the correct syntax when entering the value of the record.
- If you want to modify the field type of an unknown record, you have to delete the field type and then add it again.
- If you create an unknown record of a specific type and later on the record type is supported by NIOS, the record continues to exist as an unknown record. You will need to migrate the record to the newly supported type.
- If you add an unknown record that is not supported by the Microsoft server to the zone, you may encounter issues with the MS server synchronization.
- You cannot create records of type MD and MF.

## Modifying and Deleting Unknown Records

When you modify an unknown record, you can change the information you previously entered in the **General** tab. You can also enter or edit information in the **TTL**, **Extensible Attributes** and **Permissions** tabs. For information on modifying and deleting resource records, see [Modifying, Disabling, and Deleting Host and Resource Records](#) below.

## Prohibited Records

The following record types are prohibited as part of a zone, irrespective of whether or not they are defined as Unknown records:

- Type 0: Do not allocate it for ordinary use.
- Type 41 (OPT): Pseudo type
- Types 128-255: Meta type
- Types 55555, 55556, 55557, 65432, 65433: Used internally in NIOS
- Type 65533: Private use
- Type 3 (MD) and 4 (MF): Obsolete type



## Viewing Resource Records

You can view the configured resource records by navigating to the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *zone* -> **Records** tab. Grid Manager displays the following information for each resource record in the zone:

- **Name:** The name of the record, if applicable. For host records, this field displays the canonical name of the host. For PTR record, this displays the PTR record name without the zone name.
- **Type:** The resource record type.
- **Data:** Data that the record contains. For host records, this field displays the IP address of the host. For PTR records, this displays the domain names.
- **Active users:** The number of active users for the selected resource record.
- **Comment:** Comments that were entered for the resource record.
- **Site:** Values that were entered for this pre-defined attribute.  
Note that the DNS record that is obscured by an LBDN record is indicated by a strikethrough, for example, an obscured A record appears as A Record in Grid Manager.

You can also display the following columns:

- **MS Delegation Addresses:** This column appears only if the primary server of the zone is a Microsoft server. It displays the IP addresses that are associated with an NS record.
- **TTL:** The TTL (time-to-live) value of the record.
- **Address:** The IPv4 or IPv6 address associated with the owner domain name in a reverse-mapping zone.
- **Shared:** Displays true for shared resource records. Otherwise, displays false.
- **Shared Record Group:** Displays the shared record group name of a shared record.
- **Disabled:** Indicates if the record is disabled.

You can perform the following steps:

- Modify some of the data in the table. Double click a row and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes. Note that some fields are read only.
- Add new DNS records by clicking the arrow next to the Add icon and selecting **Host**, **Record**, **Shared Record**, and then selecting the required record type.
- View the DNS Traffic Control structure for an LBDN.
  - Select the LBDN record and click the Open Visualization icon. For more information, see [Visualization for DNS Traffic Control Objects](#).
- Create a DTC server based on an existing A, AAAA, or host record by selecting a record in the table and clicking **Create DTC Server** in the Toolbar or in the record's Action menu. For more information, see [Configuring DNS Traffic Control Servers](#).
- Edit the properties of a resource record.
  - Select the resource record, and then click the Edit icon.
- Delete a resource record.
  - Select the resource record, and then click the Delete icon.
- Export the list of resource records to a .csv file.
  - Click the Export icon.
- Print the list of resource records.
  - Click the Print icon.
- Use filters and the **Goto** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Goto** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria:
  - a. In the filter section, click **Show Filter** and define filter criteria for the quick filter.
  - b. Click **Save** and complete the configuration in the Save Quick Filter dialog box.  
The appliance adds the quick filter to the quick filter drop-down list in the panel. Note that global filters are prefixed with [G], local filters with [L], and system filters with [S].

## Modifying, Disabling, and Deleting Host and Resource Records

You can modify, disable, or delete an existing host or DNS resource record. When physical repair or relocation of a network device occurs, you can disable a record instead of deleting it. When you disable a record, the NIOS appliance does not answer queries for it, nor does it include disabled records in zone transfers and zone imports. This avoids having to delete and then add the record again. When the changes to the physical device are complete, you can simply

enable the host or resource record.

To modify or disable a host or resource record, perform the following steps:

1. Use one of the following methods to retrieve the host or resource record:
  - Perform a global search.
  - Select it from a Smart Folder.
  - From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *dns\_view* -> *zone* -> *host\_record* or *resource\_record*.
2. Select the record you want to modify and click the Modify icon.
3. In the host or resource record editor, you can do the following:
  - In the **General** tab, you can change most of the information, except for the read-only fields, such as the **DNS View** and **Host Name Policy**. You can select the **Disable** checkbox to disable the record.
  - In the **TTL** tab, you can modify the TTL setting. The NIOS appliance also allows you to specify TTL settings for each record. If you do not specify a TTL for a record, the appliance applies the default TTL value of the zone to each record. For information, see [About Time To Live Settings](#).
  - In the **Extensible Attributes** tab, you can modify the attributes. For information, see [Using Extensible Attributes](#).
  - The **Permissions** tab displays if you logged in as a superuser. For information, see [About Administrative Permissions](#).
4. Save the configuration and click **Restart** if it appears at the top of the screen.

When you delete host and resource records, Grid Manager moves them to the Recycle Bin. You can use the Recycle Bin to store deleted DNS configuration objects and selectively restore objects to the active configuration at a later time. You can also permanently remove the objects from the Recycle Bin.



#### Note

You cannot delete automatically-generated records, such as NS records and SOA records.

To delete host and resource record, perform the following steps:

1. Perform a global search to retrieve the record you want to delete.  
Or  
From the **Data Management** tab, select the **DNS** tab, click the **Zones** tab-> *dns\_view* -> *zone* -> *host\_record* or *resource\_record*.
2. Select the record and click the Delete icon.
3. In the *Delete Confirmation* dialog box, select **Yes** to delete or **No** to cancel.
4. Optionally, if the **Enable PTR record removal for A/AAAA records** option is selected and if you try to delete an A or AAAA record, the appliance displays the *Delete Confirmation (A or AAAA Record)* dialog box to confirm whether you want to remove the corresponding PTR record that was automatically generated while creating the A or AAAA record. In the *Delete Confirmation* dialog box, select the **Remove associated PTR resource record(s)** checkbox and click **Yes** to delete the associated PTR record or click **No** to cancel. For information about enabling this option, see [Deleting PTR Records associated with A or AAAA Records](#).  
Or  
You can also schedule the deletion for a later time. Click **Schedule Deletion** and in the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Deletions](#).

## Configuring Shared Record Groups

A shared record group is a set of resource records that you can add to multiple zones. You can create resource records in a group and share the group among multiple zones. The zones handle the shared resource records as any other resource record. You can include the following types of DNS resource records in a shared record group: A, SRV, MX, AAAA, CNAME, and TXT.

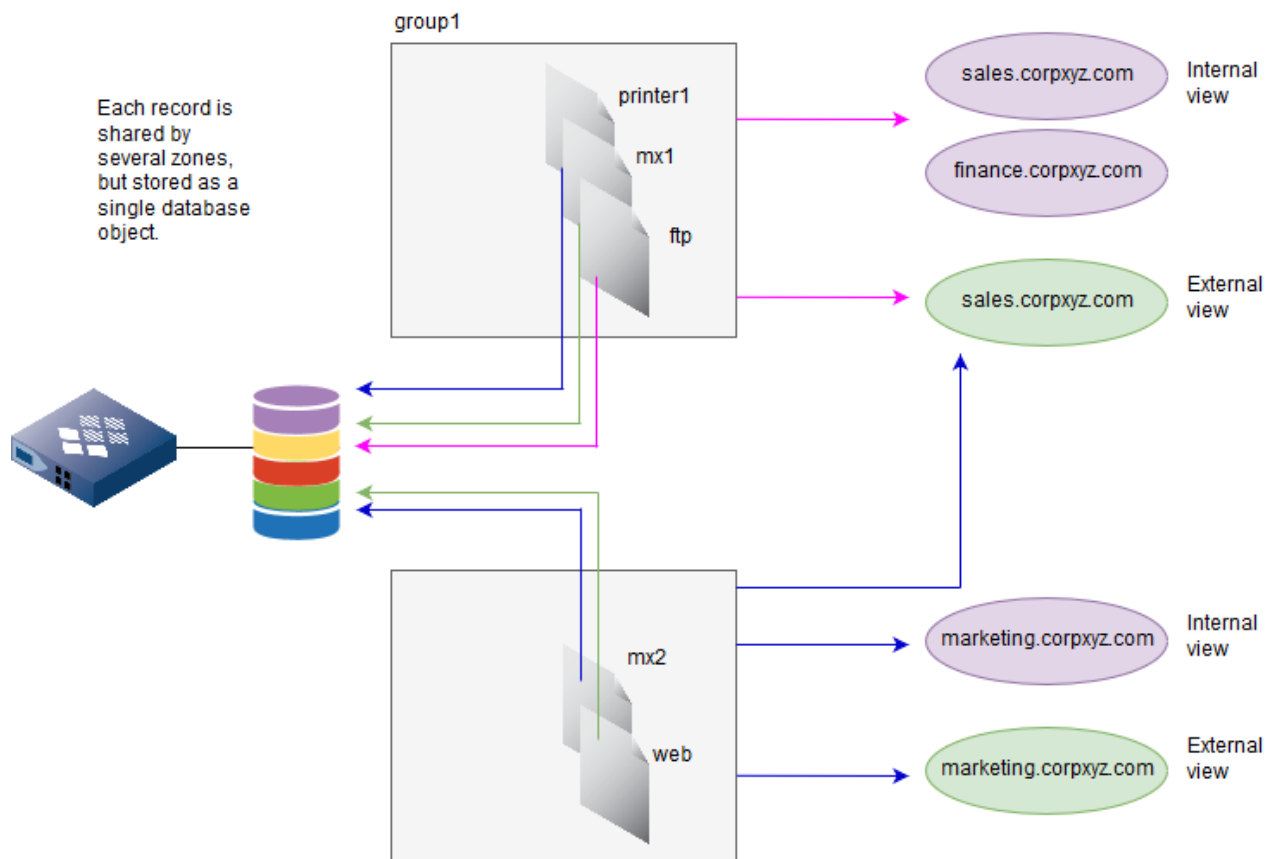
Using shared record groups simplifies and expedites the administration of resource records. When you create or update a shared record, the appliance automatically updates it in all associated zones. In addition, shared resource records reduce the object count in the NIOS database; instead of the creating the same record in multiple zones, you can use only one shared record. For example, for 10 zones and 500 records per zone, the object count decreases from 5278

objects to 781 objects.

The figure *Creating Shared Records* below shows an example of how to create and use shared records.

In this example, there are two shared record groups. One group—group1— contains the A records ftp and printer1 and the MX record mx1, and the other group—group2—contains the A record web and the MX record mx2. The resource records in group1 are shared with the internal view zones sales.corpxyz.com and finance.corpxyz.com and the external view zone sales.corpxyz.com. The resource records in group2 are shared with the internal view zone marketing.corpxyz.com and the external view zones sales.corpxyz.com and marketing.corpxyz.com.

### *Creating Shared Records*



### **Shared Records Guidelines**

The following are guidelines for using shared records:

- You can include multiple shared A, AAAA, CNAME, SRV, MX and TXT resource records in a group. You cannot include NS, DNAME, PTR, host and bulk host records.
- You can add shared records to authoritative zones only. You cannot add shared records to forward zones, stub zones, or reverse mapping zones.
- Zones that contain shared records can also contain regular DNS records (not shared).
- When you change or delete a shared resource record, it changes the canonical source of the shared record and impacts all the zones that contain the record.
- You cannot copy shared records from a zone.
- You do not need to restart the appliance when you create, delete, or modify shared records.

### **Configuring Shared Record Groups**

Before you can create shared resource records, you must first create the group to which they belong. The shared record group serves as a container for the shared resource records. The following are the tasks to configure a shared record group:

1. Create a shared record group and associate it with the appropriate zones. See [Creating a Shared Record Group](#) below.
2. Create shared A, CNAME, SRV, MX, AAAA, and TXT resource records, and add them into the shared record group. See [Managing Shared Resource Records](#) below.

## Creating a Shared Record Group

When you create a shared record group, the only requirement is that you give it a name. You can associate it with one or multiple zones when you first create the group or at a later time, by editing the shared record group. You can associate a shared record group with authoritative zones only. Associating the shared record group with a zone adds the shared records to the zone. The zone handles the shared records like any other resource records.

To create a shared record group:

1. From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab, and then click the Add icon.
2. In the *Shared Record Group* wizard, specify the following:
  - **Name:** Enter the name of the shared record group. It can be up to 64 characters long and can contain any combination of printable characters. You can change the shared record group name even after you create the group. It does not impact the shared records in the group.
  - **Hostname Policy:** Click **Override** to supersede the hostname restriction policy set at the zone level or click **Inherit** to use the zone policy. This sets the hostname policy for the shared records in the group. See [Specifying Hostname Policies](#).
  - **Comment:** Optionally, enter additional information about the shared record group.
3. Click **Next** to associate the shared record group with at least one zone.
4. Click the Add icon in the Associated Zones panel.
5. In the *Zone Selector* dialog box, select a zone by clicking the zone name. You can add multiple zones.
6. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).
7. Save the configuration.

## Viewing Shared Record Groups

You can view the configured shared record groups by navigating to the **Data Management** tab -> **DNS** tab -> **Shared Record Groups** tab. Grid Manager displays the following information about each shared record group:

- **Name:** The shared record group name.
- **Comment:** Comments that were entered for the shared record group.
- **Site:** Values that were entered for this pre-defined attribute.

You can do the following:

- List the shared resource records and associated zones in a shared record group.
- Click a shared record group name.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).
- Edit the properties of a shared record group.
  - Click the checkbox beside a shared record group, and then click the Edit icon.
- Delete a shared record group.
  - Click the checkbox beside a shared record group, and then click the Delete icon. Note that you must remove the zone associations in a shared record group before you delete it.
- Export the list of shared record groups to a .csv file.
  - Click the Export icon.
- Print the list of shared record groups.
  - Click the Print icon.

## Modifying a Shared Record Group

When you edit a shared record group, you can do the following:

1. Perform a global search to retrieve the shared record group you want to modify.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared\_record\_group* checkbox, and then click the Edit icon.
2. The *Shared Record Group* editor contains the following tabs from which you can modify information:
  - **General**: You can change any of the information you entered when you created it, including its name. Changing the shared record group name does not impact the shared resource records in it.
  - **Extensible Attributes**: You can modify the attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions**: This tab is displayed if you logged in as a superuser. For information, see [About Administrative Permissions](#).
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Deleting Shared Record Groups

Before you delete a shared record group, you must remove the zone associations in the group; otherwise, an error message appears when you delete. For information, see [Deleting Associated Zones](#) below.

To delete a shared record group:

1. Perform a global search to retrieve the shared record group you want to modify.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared\_record\_group* checkbox, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.

Grid Manager moves the shared record group to the Recycle Bin, if enabled. Use the Recycle Bin feature to recover a deleted shared record group and retrieve the deleted zones. For information, see [Using the Recycle Bin](#).

## Managing Shared Resource Records

You can create shared A, AAAA, CNAME, MX, SRV and TXT records. These resource records are similar to the non-shared resource records. The DNS server uses them to respond to queries in the same way as any other resource record. A shared resource record can belong to only one shared record group. This section describes how to add shared resource records to a group and how to modify and delete them. It includes the following sections:

- [Creating Shared Records](#)
- [Viewing Shared Records](#)
- [Modifying Shared Records](#)
- [Deleting Shared Records](#)

NIOS allows superusers to add shared A and AAAA records with a blank name. Limited-access users must have read/write permission to **Adding a blank A/AAAA record** to add shared A and AAAA records with a blank name. You can assign global permission for specific admin groups and roles to allow limited-access users to add shared A and AAAA records with a blank name. For more information, see [Administrative Permissions for Adding Blank A or AAAA Records](#).

## Creating Shared Records

After you create a shared record group, you can create its resource records. To create a shared A, AAAA, CNAME, MX, SRV or TXT record and add it to a group:

1. From the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **Add** -> **Shared Record**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared\_record\_group*. Expand the **Shared Records** tab and click the Add icon.
2. Select one of the following:
  - **Shared A Record**

- **Shared AAAA Record**
  - **Shared CNAME Record**
  - **Shared MX Record**
  - **Shared SRV Record**
  - **Shared TXT Record**
3. Enter information in the *Shared Record* wizard. See the online Help or the following for information about each resource record:
    - For information about A records, see [Managing A Records](#).
    - For information about AAAA records, see [Managing AAAA Records](#).
    - For information about CNAME records, see [Managing CNAME Records](#).
    - For information about MX records, see [Managing MX Records](#).
    - For information about SRV records, see [Managing SRV Records](#).
    - For information about TXT records, see [Managing TXT Records](#).
  4. Save the configuration, or click **Next** to define extensible attributes for the shared record. For information, see [Using Extensible Attributes](#).
  5. Click **Restart** if it appears at the top of the screen.

### Viewing Shared Records

You can view the shared records in a group and in a zone. To edit the shared record properties, click the shared record name and select the Edit icon.

To view the shared records in a group:

- From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared\_record\_group* -> **Shared Records** tab.

To view the shared records in a zone:

- From the **Data Management** tab, select the **DNS** tab -> **Zones** tab and select a zone. Grid Manager lists the following information about each shared record by default:
  - **Name:** The shared record name.
  - **Type:** Indicates the type of resource record, such as A, AAAA, CNAME, MX, SRV or TXT records. Shared records are identified as **(Shared)**.
  - **Data:** The data the shared resource record provides.
  - **Comment:** Comments that were entered in the resource record.
  - **Site:** Displays values that were entered for this pre-defined attribute.

You can display the following additional columns:

- **TTL:** The TTL value of the shared resource record.
- **Disabled:** Indicates whether the record is disabled.

You can do the following:

- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).
- Edit the properties of a shared resource record.
  - Select the shared resource record, and then click the Edit icon.
- Delete a shared resource record.
  - Select the shared resource record, and then click the Delete icon.
- Export the list of shared resource records to a .csv file.
  - Click the Export icon.
- Print the list of shared resource records.
  - Click the Print icon.

## Modifying Shared Records

You can modify, disable, or delete any shared record. When physical repair or relocation of a network device occurs, you can disable a record instead of deleting it. This alleviates having to delete, and then add the shared record again. When the changes to the physical device are complete, you can simply enable the shared record.

To modify or disable a shared record:

1. Perform a global search to retrieve the host or resource record you want to modify.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared\_record\_group* -> **Shared Records** tab.
2. Select the shared record you want to modify and click the Edit icon.
3. The *Shared Records* editor contains the following tabs from which you can modify information:
  - **General**: You can change most of the information, except for the read-only fields, such as the Host Name Policy. You can also select the **Disable** checkbox to disable the record.
  - **TTL**: You can modify the TTL setting. For information, see [About Time To Live Settings](#).
  - **Extensible Attributes**: You can modify the attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions**: This tab displays if you logged in as a superuser. For information, see [About Administrative Permissions](#).
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Deleting Shared Records

To delete shared resource records:

1. Perform a global search to retrieve the record you want to delete.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared\_record\_group* -> **Shared Records** tab.
2. Select the shared record you want to delete and click the Delete icon.
3. When the confirmation dialog box displays, select **Yes**.

Grid Manager moves the shared records to the Recycle Bin, from which you can restore or permanently delete the records.

## Managing Associated Zones

Typically, you associate a zone with a shared record group when you create the group. You can also add an associated zone to a shared record group after you create the group.

## Creating Associated Zones

To associate a zone with a share record group:

1. From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared\_record\_group* -> **Associated Zones** tab, and then click the Add icon.
2. In the *Zone Selector* dialog box, select a zone by clicking the zone name.

The appliance adds the zone to the **Associated Zones** tab.

## Viewing Associated Zones

To view the associated zones in a shared record group:

- From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared\_record\_group* -> **Associated Zones** tab.

Grid Manager lists the following information about each associated zone by default:

- **Zone**: The zone associated with the shared record group.



- **DNS View:** The DNS view to which the zones belong.
- **Network View:** The network view associated with the DNS view.
- **Comment:** Comments that were entered for the shared record group. You can do the following:
  - Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
  - Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Associate another zone with the shared record group.
  - Click the Add icon and select a zone.
- Delete an associated zone.
  - Select the zone, and then click the Delete icon.
- Export the list of associated zones to a .csv file.
  - Click the Export icon.
- Print the list of shared associated zones.
  - Click the Print icon.

### Deleting Associated Zones

To delete an associated zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab -> *shared\_record\_group* -> **Associated Zones** tab.
2. Select the associated zone and click the Delete icon.
3. When the confirmation dialog box displays, select **Yes**.

Grid Manager removes the zone from the shared record group.

### Configuration Example: Configuring Shared Records

The following example shows you how to configure shared records. In this example, you do the following:

- Create a shared record group: **group1**.
  - Associate it with three zones: **eng.com**, **sales.com**, and **marketing.com**.
  - Create an A record **www** and an MX record **mx1**.
1. Create a shared record group called **group1** and associate it with **eng.com**, **sales.com**, and **marketing.com**.
    - a. From the **Data Management** tab, select the **DNS** tab -> **Shared Record Groups** tab, and then click the Add icon.
    - b. In the first step of the *Shared Record Group* wizard, specify the following **Name:** Enter **group1**.
    - c. Click **Next**.
    - d. Click the Add icon in the Associated Zones panel.
    - e. Select **eng.com** from the list of zones and click the select icon. Do the same for the **sales.com**, and **marketing.com** zones.
    - f. Save the configuration and click **Restart** if it appears at the top of the screen.
  2. Add an A record **www** to **group1**.
    - a. Expand the Toolbar and click **Add** -> **Shared Record** > **Shared A Record**.
    - b. In the *Shared A Record* wizard, specify the following:
      - Name:** Enter **www**.
      - Shared Record Group:** Select **group1** from the drop-down list.
      - IP Address:** Enter the IP address **10.9.1.1**.
    - c. Save the configuration and click **Restart** if it appears at the top of the screen.
  3. Add an MX record **mx1** into **group1**.
    - a. Expand the Toolbar and click **Add** -> **Shared Record** > **Shared MX Record**.
    - b. In the *Shared MX Record* wizard, specify the following:
      - Mail Destination:** Enter **mx1**.
      - Shared Record Group:** Select **group1** from the drop-down list.
      - Mail Exchanger:** Enter **www.infoblox.com**.
      - Preference:** Enter **10**.
      - Comment:** Enter **mail exchanger record for shared record group1**.
    - c. Save the configuration and click **Restart** if it appears at the top of the screen.



## Configuring DDNS Updates

DDNS (Dynamic DNS) is a method to update DNS data (A, TXT, and PTR records) from sources such as DHCP servers and other systems that support DDNS updates, such as Microsoft Windows servers 2000, 2003, 2008, 2008 R2, 2012, 2012 R2, and 2016.

This section provides conceptual information about DDNS and explains how to configure NIOS appliances running DHCP, DHCPv6 and DNS to support DDNS updates. It contains the following main topics:

- [Understanding DDNS Updates from DHCP](#)
- [Configuring DHCP for DDNS](#)
- [Configuring DDNS Features](#)
- [About the Client FQDN Option](#)
- [Configuring DDNS Update Verification](#)
- [Configuring DNS Servers for DDNS](#)
- [Supporting Active Directory](#)
- [About GSS-TSIG](#)
- [Configuring GSS-TSIG keys](#)
- [Accepting DDNS Updates from DHCP Clients](#)
- [Accepting GSS-TSIG-Authenticated Updates](#)
- [Secure Dynamic Updates](#)

## Understanding DDNS Updates from DHCP

DHCP supports several DNS-related options (such as options 12, 15, and 81 for IPv4, and options 23, 24, and 39 for IPv6). With DDNS (Dynamic DNS) updates, a DHCP server or client can use the information in these options to inform a DNS server of dynamic domain name-to-IP address assignments.




### Note

DDNS updates is not supported by IPv6-only appliances.

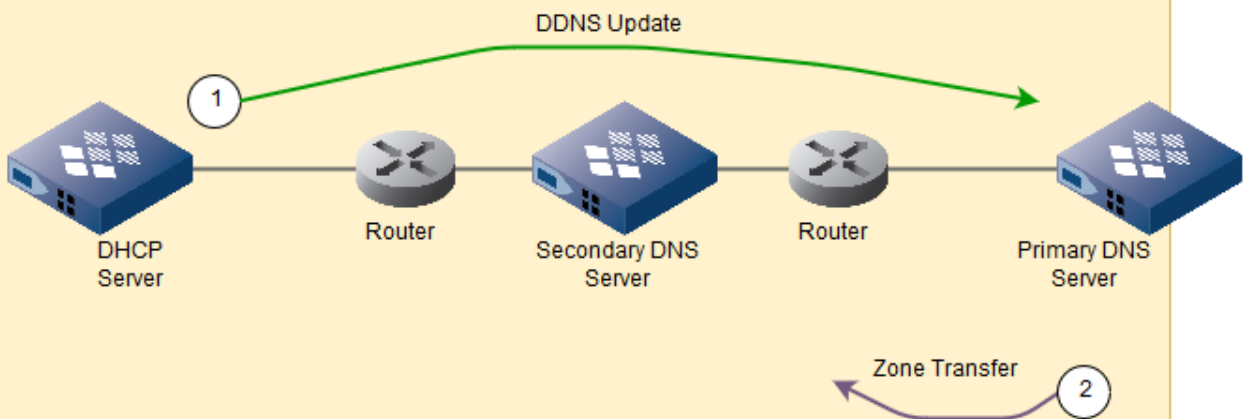
To set up one or more NIOS appliances for DDNS updates originating from DHCP, you must configure at least one DHCP server and one DNS server. These servers might be on the same appliance or on separate appliances. Three possible arrangements for a DHCP server to update a DNS server are shown in the figure below.

*Relationship of DHCP and DNS Servers for DDNS Updates*

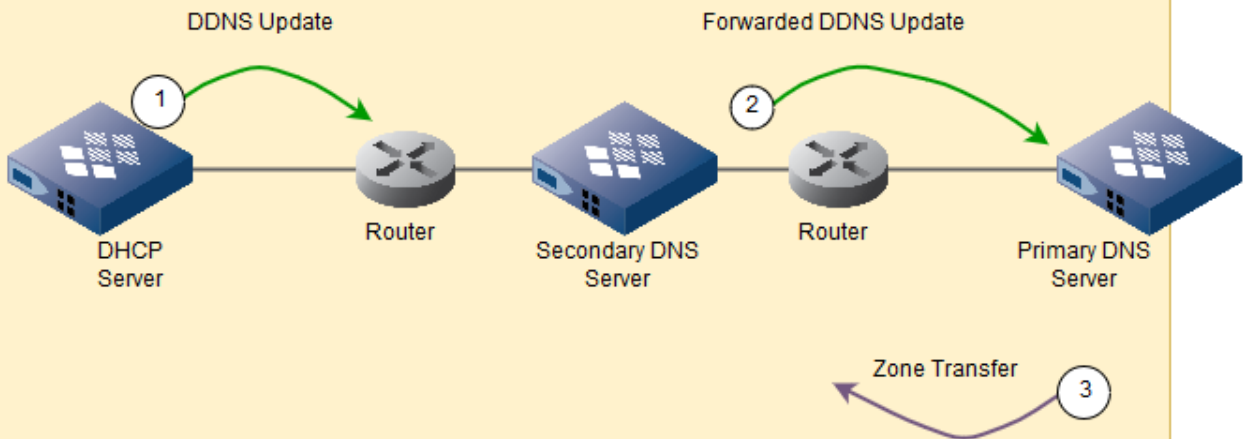
DDNS when the DHCP server and primary DNS server are on the same NIOS appliance.



DDNS when the DHCP server and primary DNS server are on the different NIOS appliances and the DHCP server updates the primary DNS server.



DDNS when the DHCP server and primary DNS server are on the different appliances and the DHCP server updates a secondary DNS server.



Here is a closer look at one setup for performing DDNS updates from a DHCP server (the steps relate to figure [DDNS Update from a DHCP Server](#) below).

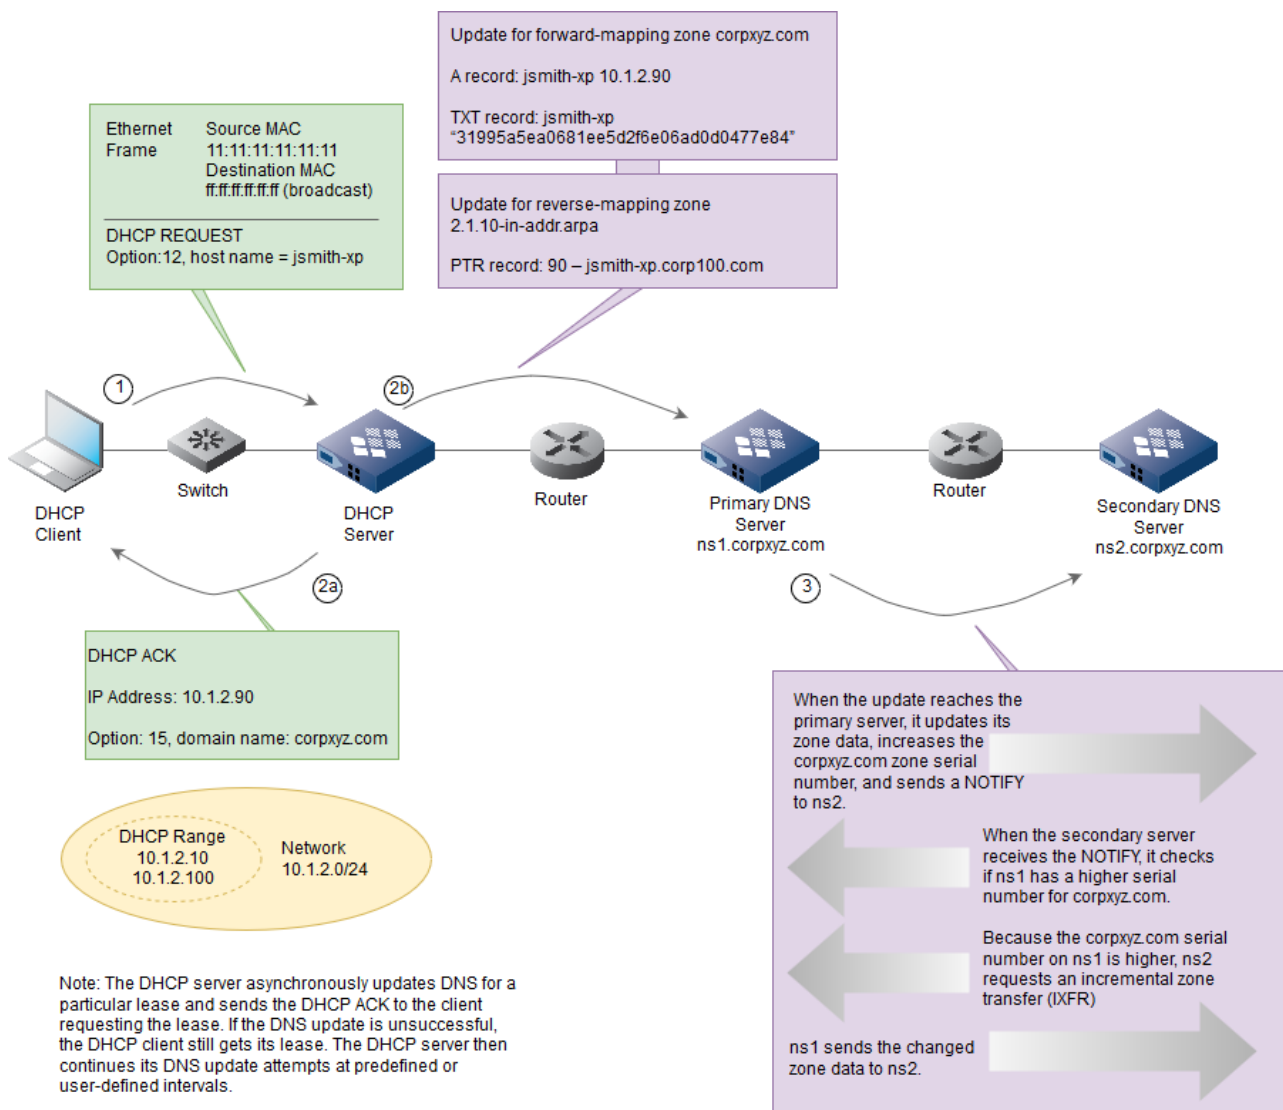
1. When an IPv4 DHCP client requests an IP address, the client sends its host name (DHCP option 12). The client also includes its MAC address in the ethernet frame header.
2.
  - a. When the DHCP server responds with an IP address, it usually provides a domain name (DHCP option 15). The combined host name (from the client) and domain name (from the server) form an FQDN (fully qualified domain name), which the NIOS appliance associates with the IP address in the DHCP lease.
  - b. The DHCP server sends the A, TXT, and PTR records of the DHCP client to the primary DNS server to update its resource records with the dynamically associated FQDN + IP address.
3. The primary DNS server notifies its secondary servers of a change. The secondary servers confirm the need for a zone transfer, and the primary server sends the updated zone data to the secondary server, completing the update.



Note

For information about zone transfers, see [Enabling Zone Transfers](#).

*DDNS Update from a DHCP Server*



To enable a DHCP server to send DDNS updates to a DNS server, you must configure both servers to support the updates. First, configure the DHCP server to do the following:

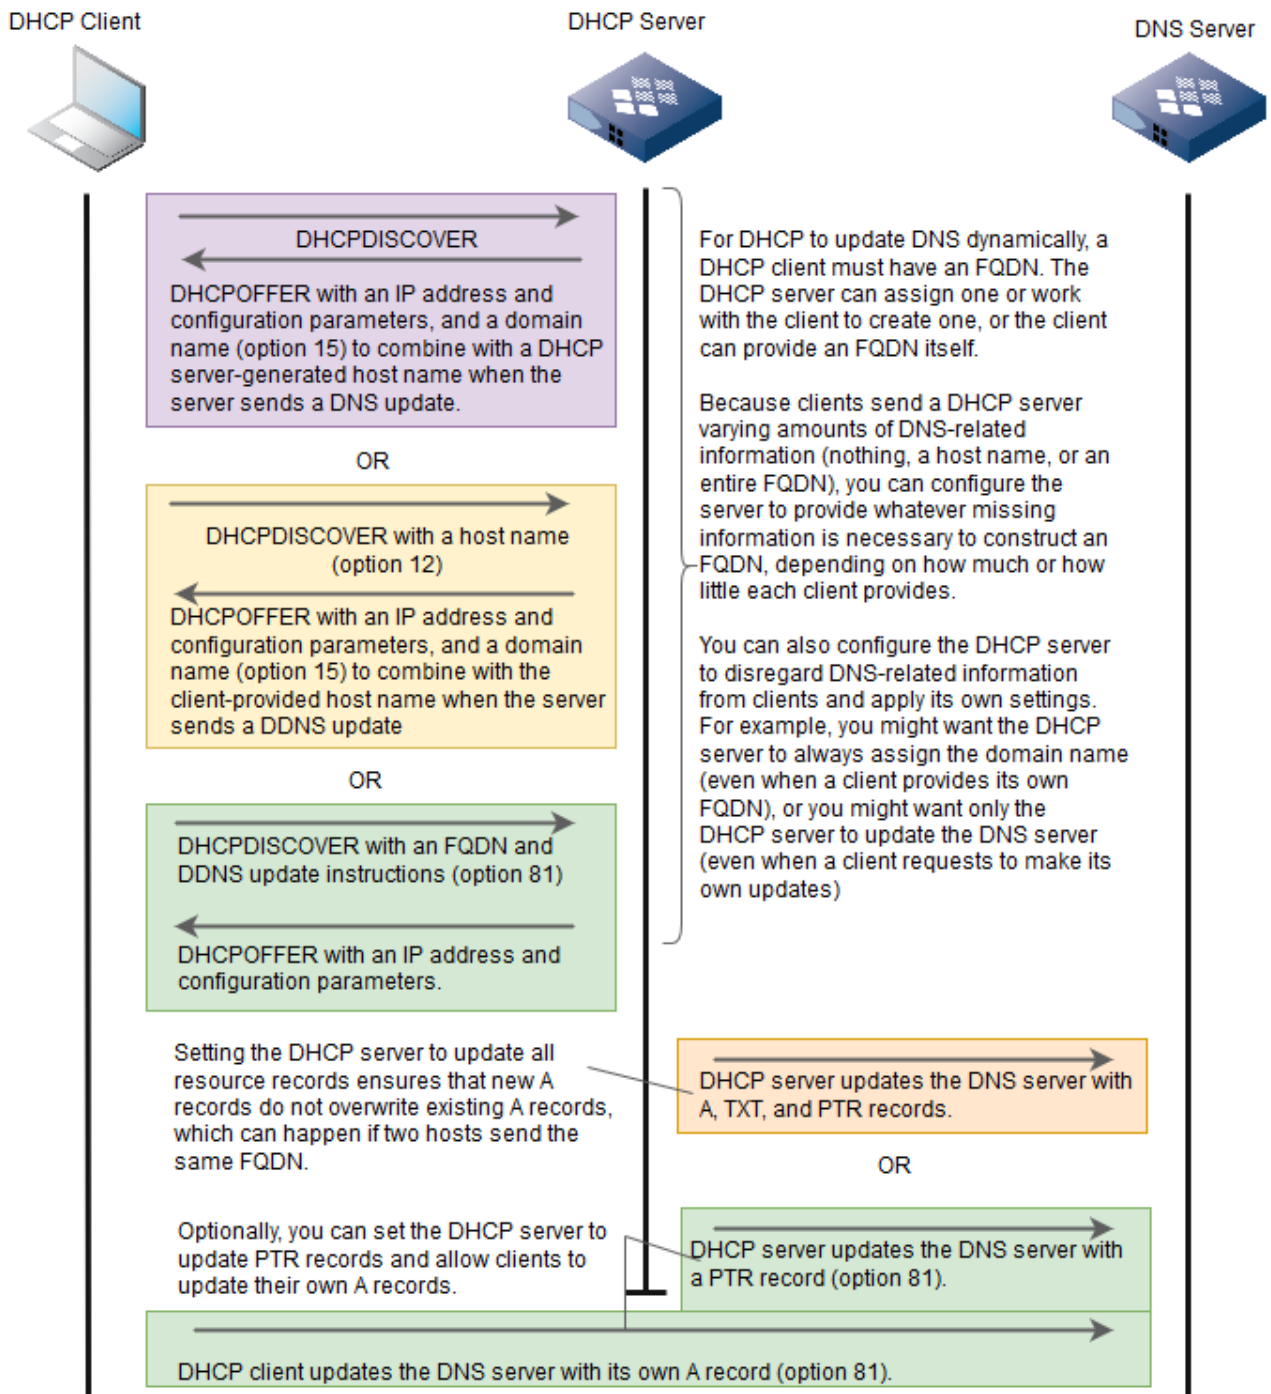
- Provide what is needed to create an FQDN: add a server-generated host name to a server-provided domain name, add a server-provided domain name to a client-supplied host name, or permit the client to provide its own FQDN
- Send updates to a DNS server

Then, configure the following on the DNS server:

- Accept updates from the DHCP server, a secondary DNS server, or a DHCP client
- If the DHCP server sends updates to a secondary DNS server, configure the secondary server to forward updates to the primary DNS server

When setting up DDNS, you can determine the amount of information that DHCP clients provide to a DHCP server — and vice versa — and where the DDNS updates originate. A summary of these options for IPv4 is shown in the following figure. It is similar for IPv6, except that the DHCP client and server exchange Request and Reply messages, AAAA records are updated instead of A records, and the FQDN option is option 39.

#### *DHCP Clients and Server Providing DNS Information and Updates*



You can configure the DHCP and DNS settings for DDNS at the Grid level, member level, and network and zone level. By applying the inheritance model in the NIOS appliance, settings made at the Grid level apply to all members in the Grid. Settings you make at the member level apply to all networks and zones configured on that member. Settings made at the network and zone level apply specifically to just that network and zone. When configuring independent appliances (that is, appliances that are not in a Grid), do not use the member-level settings. Instead, configure DDNS updates at the Grid level to apply to all zones and, if necessary, override the Grid-level settings on a per zone basis.

## Configuring DHCP for DDNS

Before a DHCP server can update DNS, the DHCP server needs to have an FQDN-to-IP address mapping. When a DHCP IPv4 client requests an IP address, it typically includes its host name in option 12 of the DHCPDISCOVER packet, and an IPv6 client includes its hostname in the Request packet. You can configure the NIOS appliance to include a domain name in option 15 of the IPv4 DHCPOFFER packet or in the IPv6 Reply packet. You specify this domain name in the **IPv4 DHCP Options** -> **Basic** and **IPv6 DHCP Options** -> **Basic** tabs of the *Grid DHCP Properties* editor, *Member DHCP Configuration* editor, and the *Network* editor. For IPv4 clients you can also specify a domain name in the *DHCP Range* and *Fixed Address* editors.

Then, you can enable the DHCP server to send DDNS updates for IPv4 and IPv6 clients, as described in [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#) below. After you enable the DHCP server to send DDNS updates, you can do the following:

- Configure the DHCP server to send DDNS updates to DNS servers in the Grid. For information, see [Sending Updates to DNS Servers in the Grid](#) below.
- Configure the DHCP server to send DDNS updates to external DNS servers. For information, see [Configuring DDNS Features](#).
- Configure certain DDNS features. For information, see [Configuring DDNS Features](#).
- Enable support for the FQDN option for IPv4 and IPv6 clients, and configure how the DHCP server updates DNS. For information, see [Enabling FQDN Option Support](#).



### Note

Whether you deploy NIOS appliance in a Grid or independently, they send updates to UDP port 53. Grid members do not send updates through a VPN tunnel; however, Grid members do authenticate updates between each other using TSIG (transaction signatures) based on an internal TSIG key.

## Enabling DDNS for IPv4 and IPv6 DHCP Clients

You can enable the DHCP server to send DDNS updates for IPv4 clients at the Grid, member, shared network, network, address range, DHCP template, fixed address, and roaming host levels, and for IPv6 clients at the Grid, member, network, shared network, network template and roaming host levels.

You can specify a different domain name that the appliance uses specifically for DDNS updates. The appliance combines the hostname from the client and the domain name you specify to create the FQDN that it uses to update DNS. For IPv4 clients, you can specify the DDNS domain name at the network, network template, range, and range template levels. For IPv6 clients, you can specify the DDNS domain name at the Grid, member, network, shared network, and network template levels. You can also use the name of a roaming host record as the name of the client for DDNS updates, as described in [Setting Properties for Roaming Hosts](#).

To enable DDNS and specify a DDNS domain name:

1. **Grid:** From the **DataManagement** tab, select the **DHCP** tab, expand the Toolbar and click **GridDHCPProperties**.  
**Member:** From the **DataManagement** tab, select the **DHCP** tab and click the **Members** tab -> **Members** -> *member* checkbox -> Edit icon.  
**Network:** From the **DataManagement** tab, select the **DHCP** tab and click the **Networks** tab -> **Networks** -> *network* checkbox -> Edit icon.  
**NetworkContainer:** From the **DataManagement** tab, select the **IPAM** tab -> *network\_container* checkbox, and then click the Edit icon.  
**NetworkTemplate:** From the **DataManagement** tab, select the **DHCP** tab and click the **Templates** tab -> *DHCP\_template* checkbox -> Edit icon.  
**RoamingHost:** From the **DataManagement** tab, select the **DHCP** tab and click the **Networks** tab -> **RoamingHosts** -> *roaming\_host* -> Edit icon.  
For IPv4 clients only: **IPv4AddressRange:** From the **DataManagement** tab, select the **DHCP** tab and click the **Networks** tab -> **Networks** -> *network* -> *addr\_range* checkbox -> Edit icon.  
**IPv4FixedAddress:** From the **DataManagement** tab, select the **DHCP** tab and click the **Networks** tab -> **Networks** -> *network* -> *ip\_addr* checkbox -> Edit icon.

**IPv4AddressRange/FixedAddressTemplate:** From the **DataManagement** tab, select the **DHCP** tab and click the **Templates** tab -> *DHCP\_template* checkbox -> Edit icon.

To override an inherited property, click **Override** next to it and complete the appropriate fields.

2. In the **IPv4DDNS** -> **Basic** tab or the **IPv6DDNS** -> **Basic** tab, complete the following:
  - **Enable DDNS Updates:** Select this checkbox to enable DDNS updates.  
When setting properties for DHCP objects other than the Grid, you must click **Override** and select **Enable DDNS updates** for the DDNS settings to take effect. When turning on DDNS updates, first verify if [Option 81](#) has been enabled and whether DNS is being updated. If DNS is being updated, even if the DNS zone targets are not in the Grid, select Option 81 support and the correct sub option. For more information, see [Enabling FQDN Option Support](#).  
When the **Enable DDNS Updates** checkbox is not selected, the default behaviour is to allow the client to update DNS.  
When the **Enable DDNS Updates** checkbox is selected, the default behaviour is to prevent DDNS updates from the client.  
In a dual mode Grid, if IPv6 DDNS updates is enabled at the Grid level, then when you join an IPv6 Grid member to the Grid, IPv6 DDNS updates is automatically disabled for the Grid member.
  - **DDNS domain name:** Specify the domain name of the network that the appliance uses to update DNS. For IPv4 clients, you can specify this at the network, network template, range, and range template levels. For IPv6 clients, you can specify this at the Grid, member, network, shared network, and network template levels.
  - **DDNS Update TTL:** You can set the TTL used for A or AAAA and PTR records updated by the DHCP server. The default is shown as zero. If you do not enter a value here, the appliance by default sets the TTL to half of the DHCP lease time with a maximum of 3600 seconds. For example, a lease time of 1800 seconds results in a TTL of 900 seconds, and a lease time of 86400 seconds results in a TTL of 3600 seconds. For information about how to set the lease time, see [Defining Lease Times](#).
  - **DDNS Update Method:** Select the method used by the DHCP server to send DDNS updates. You can select either **Interim** or **Standard** from the drop-down list. The default is **Interim**. When you select **Interim**, TXT record will be created for DDNS updates and when you select **Standard**, DHCID record will be created for DDNS updates. But in the **IPv4 DDNS** -> **Advanced** tab or the **IPv6 DDNS** -> **Advanced** tab, if you have selected **No TXT Record** mode for the DHCP server to use when handling DNS updates, then TXT record or DHCID record is not created for DDNS updates.  
If you change the DDNS update method from **Interim** to **Standard** or vice versa, then the DHCP server changes the DHCID type used from TXT record to DHCID record or vice versa as the leases are renewed.  
This is supported for clients that acquire both IPv4 and IPv6 leases. Infoblox recommends you to configure different DDNS update method for IPV4 leases and IPV6 leases, **Interim** for IPv4 lease and **Standard** for IPv6 lease.
  - **Update DNS on DHCP Lease Renewal:** Select this checkbox to enable the appliance to update DNS when a DHCP lease is renewed.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Sending Updates to DNS Servers

The DHCP server can send DDNS updates to DNS servers in the same Grid and to external DNS servers. When you enable the appliance to send updates to Grid members, you must specify the DNS view to be updated. If a network view has multiple DNS views, you can select only one DNS view for DDNS updates. For information about DNS views, see [About DNS Views](#).

When you enable DDNS updates for a Grid, member, shared network, network, address range, DHCP template, fixed address, or roaming host, the DHCP server sends updates to authoritative zones using the domain name (as DHCP option 15) you define in the DHCP properties. You can also define forward-mapping zones that receive DDNS updates for DHCP clients that use option 81 to define the domain name. For information, see [About the Client FQDN Option](#). To allow DDNS updates for clients using option 81, you must first enable the support for option 81. For information, see [Configuring DDNS Features](#).

For DNS zones that have multiple primary servers, you can define a primary name server to be used as the default primary server when performing DDNS updates from the appliance. Note that you cannot configure an external primary as the default primary. For more information, see [Defining the Default Primary for DDNS Updates to Zones with Multiple Primaries](#) below.



## Sending Updates to DNS Servers in the Grid

You must specify the DNS view to be updated for each network view.

To configure the DHCP server to send updates to DNS servers in the same Grid:

1. If there are multiple network views in the Grid, select a network view.
2. From the **Data Management** tab, select the **DHCP** tab, and then click **Configure DDNS** from the Toolbar.
3. In the *DDNS Properties* editor, complete the following:
  - **DNS View:** If a network view has more than one DNS view, this field lists the associated DNS views. From the drop-down list, select the DNS view to which the DHCP server sends DDNS updates. Otherwise, the appliance uses the default DNS view.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

The appliance sends DDNS updates to the appropriate zones in the selected DNS view. Note that you cannot delete a DNS view that has been selected for DDNS updates. By default, the DHCP server sends DDNS updates to zones using the domain name that you define for DHCP objects, such as networks and DHCP ranges.

## Sending Updates for Zones on an External Name Server

The DHCP server can send dynamic updates to an external name server that you specify. For each network view, you can specify the zone to be updated and the IP address of the primary name server for that zone. You can add information for a forward and reverse zone. The DHCP server updates the A record in the forward zone and the PTR record in the reverse zone.

You can also use TSIG (transaction signatures) or GSS-TSIG to secure communications between the servers. TSIG uses the MD5 (Message Digest 5) algorithm and a shared secret key to create an HMAC (hashed message authentication code)—sometimes called a digital fingerprint—of each update. Both the DHCP server sending the update and the DNS server receiving it must share the same secret key. Also, it is important that the time stamps on the TSIG-authenticated updates and update responses be synchronized, or the participants reject them. Therefore, use an NTP server to set the time on all systems involved in TSIG authentication operations.

Note that only a superuser can configure DDNS. To configure DDNS, a limited-access admin must contact a superuser.

To send updates to a DNS server that is external to your Grid:

1. If there are multiple network views in the Grid, select a network view.
2. From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Configure DDNS**.
3. In the **DDNS Updates to External Zones** section of the *DDNS Properties* editor, click the Add icon. Complete the following fields in the Add External DDNS Zone panel, and then click **Add**:
  - **Zone Name:** Enter the FQDN of a valid forward-mapping or reverse-mapping zone to which the DHCP server sends the updates. Do not enter the zone name in CIDR format. To specify a zone name in IDN, manually convert IDN to punycode and use the punycode representation.
  - **DNS Server Address:** Enter the IP address of the primary name server for that zone.
  - **Security:** Select one of the following security methods:
    - **None:** Select this to use unsecured DDNS updates. This is the default.
    - **TSIG:** Select this to use the standards-based TSIG key that uses the one-way hash function MD5 to secure transfers between name servers. You can either specify an existing key or generate a new key.  
To specify an existing key, complete the following:  
**Key Name:** Enter the TSIG key name. The key name entered here must match the TSIG key name on the external name server.  
**Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.  
**Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **GenerateKeyData** drop down list, and then click **Generate Key Data** to create a new key.
    - **GSS-TSIG:** For information about using GSS-TSIG, see [About GSS-TSIG](#).
4. Save the configuration and click **Restart** if it appears at the top of the screen.



## Defining the Default Primary for DDNS Updates to Zones with Multiple Primaries

If you have configured multiple primary servers for an authoritative zone, you can define the default primary that the appliance uses to perform DDNS updates for the zone. Note that you can configure a Grid primary, but not an external primary, as the default primary. If you do not configure a default primary, the Grid Master becomes the default primary for the zones that it serves. Otherwise, the appliance selects a primary server that serves the zone as the default primary. For external zones that have multiple primaries, the first external primary server becomes the default primary.

Configuring a default primary for DDNS updates is useful when you have DHCP members that span across different locations. Performing DDNS updates becomes more efficient when you configure a default primary that is close in proximity to the DHCP member. For example, zone corpxyz.com has two primaries (usa.corpxyz.com and japan.corpxyz.com) serving two locations (USA and Japan). Service performance is faster when you select usa.corpxyz.com as the default primary for DDNS updates in the USA region and japan.corpxyz.com as the default primary for the Japan region.

When you configure a preferred or default primary server for DDNS updates to a zone that has multiple primaries, ensure that the following are in place:

- The zone that you select contains multiple primary servers.
- The primary server has DNS service enabled and is authoritative for the zone.
- The appliance has DHCP service enabled.



### Note

You can define the default primary for the Grid and override the setting at the member level, and you must restart service for the configuration to take effect. Primary selection is performed at service restart, not at runtime.

To define the default primary:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Configure DDNS**. In the DDNS Properties editor, scroll down to the **Master Preferences for DDNS Updates to Multi-master DNS Zones** section.  
**Member:** From the **Data Management** tab, select the **DHCP** tab and click the **Members** tab -> *member* checkbox -> Edit icon. In the *Member DHCP Properties* editor, select the **Multi-master DDNS** tab.
2. In the **Master Preferences for DDNS Updates to Multi-master DNS Zones** section, click the Add icon and select one of the following from the list:
  - **Grid Zone:** In the Add Grid Zone panel, complete the following:
    - **Zone:** Click **Select Zone** to select the zone that has multiple primaries. Note that when configuring for the Grid, only zones that belong to the selected DNS view are displayed in the *Zone Selector* dialog box. For a Grid member however, if it is associated with a network view or if you have defined a DNS view for DDNS updates for that network view, only zones that belong to the DNS view configured for the associated network view are displayed.
    - **DNS View:** Displays the DNS view to which this zone belongs.
    - **DNS Primary:** From the drop-down list, select the primary name server you want the appliance to use when performing DDNS updates. Note that the list displays only primary servers that are defined for the selected zone.
  - **Default Primary:** In the Add Default Primary panel, complete the following:
    - **DNS Primary:** Click **Select** to select a primary name server from the *Member Selector* dialog. When you select a default primary, the appliance uses this name server for DDNS updates to all zones.
3. Click **Add** to add the zone and primary name server to the table, which displays the following information:
  - **DNS View:** The DNS view to which the zone belongs.
  - **Zone:** The selected zone that has multiple primaries. All zones added to the table belong to the same DNS view.
  - **DNS Primary:** The primary server to be used when performing DDNS updates from a NIOS DHCP server to the selected zone. If you have configured more than one Grid DNS primary server for DDNS updates for multi-master zones, DHCP servers use the first available DNS primary server that is configured. If the first DNS primary server is not reachable or is offline, then the DHCP servers reach for the next DNS

primary server in the preferred multi-domain DDNS list and so on. You can add upto a maximum of three DNS primary nameservers for each zone.

4. **Concatenated with the following rules defined at the Grid level:** This section appears only in the *Member DHCP Properties* editor. This table displays rules that are defined for zones with multiple primaries at the Grid level. Rules configured at the member level automatically override those configured for the Grid. Note that all rules configured for both the Grid and the member apply.

## Configuring DDNS Features

You can enable the DHCP server to support certain DDNS features for IPv4 and IPv6 clients. These features affect the behavior of the DHCP server and how it handles DDNS updates. The following sections describe the different features you can set.

### Resending DDNS Updates

You can enable the DHCP server to make repeated attempts to send DDNS updates to a DNS server. The DHCP server asynchronously updates DNS for a particular lease and sends the DHCP ACK to the client requesting the lease. If the update fails, the DHCP server still provides the lease and sends the DHCP ACK to the client. The DHCP server then continues to send the updates until it is successful or the lease of the client expires. You can change the default retry interval, which is five minutes.

You can enable this feature for the Grid and for individual Grid members.

### Generating Host Names for DDNS Updates

Some IPv4 and IPv6 clients do not send a host name with their DHCP requests. When the DHCP server receives such a request, its default behavior is to provide a lease but not update DNS. You can configure the DHCP server to generate a host name and update DNS with this host name when it receives a DHCP request that does not include a host name. It generates a name in the following format: `dhcp-ip_address`, where *ip\_address* is the IP address of the lease. For example, if this feature is enabled and the DHCP server receives a DHCP REQUEST from an IPv4 DHCP client with IP address 10.1.1.1 and no host name, the DHCP server generates the name `dhcp-10-1-1-1` and appends the domain name, if specified, for the DDNS update. Likewise, if an IPv6 client with IP address 2001:db8:a23:0:0:0:d sends a request, the DHCP server generates the name `dhcp-2001-db8-a23-0-0-0-d` and appends the domain name, if specified, for the DDNS update.

### Updating DNS for IPv4 Clients with Fixed Addresses

By default, the DHCP server does not update DNS when it allocates an IPv4 or IPv6 fixed address to a client. You can configure the DHCP server to update the A and PTR record of IPv4 clients with a fixed address. When you enable this feature and the DHCP server adds A and PTR records for a fixed address, the DHCP server never discards the records. When the lease of the client terminates, you must delete the records manually. Note that the DHCP server does not send DDNS updates for IPv6 fixed addresses and hosts. You can define fixed address settings for the Grid, Grid members, IPv4 networks, and IPv4 shared networks.

## Configuring DDNS Features

You can configure DDNS features for a Grid, its member, IPv4 and IPv6 networks and shared networks, and IPv4 DHCP address ranges. You cannot set DDNS features for IPv6 DHCP ranges. To configure DDNS features:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Properties**.  
**Member:** From the **Data Management** tab, select the **DHCP** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
**Network:** From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> **Networks** -> *network* checkbox -> Edit icon.  
**Shared Network:** From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> **Shared Networks** -> *shared\_network* checkbox -> Edit icon.

**DHCP Range:** From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> **Networks** -> *network* -> *addr\_range* checkbox -> Edit icon.

To override an inherited property, click **Override** next to it and complete the appropriate fields.

2. In the **DDNS** -> **Advanced** tab for the Grid and member, or the **DDNS** -> **Basic** tab for the network, do the following:
  - **Update Retry:** You can set this for a Grid and its members only.
    - **Retry Updates When Server Becomes Available:** Select this checkbox.
    - **Retry interval (Minutes):** You can optionally set the retry interval. The default is five minutes.
  - **Generate Hostname**
    - **Generate Hostname if not Sent by Client:** Select this checkbox to enable the DHCP server to generate a hostname and update DNS with this hostname, when the DHCP request of a client does not include a hostname.
  - **Fixed Address Updates:** You can set this for IPv4 fixed addresses only. This option is available in the **IPv4 DDNS Advanced** tab of the *Grid DHCP Properties* and *Member DHCP Properties* editors, and in the **IPv4 DDNS Basic** tab of the *IPv4 DHCP Network* and *Shared Network* editors.
    - **Update Fixed Addresses:** Select this checkbox to allow the DHCP server to send updates to DNS for IPv4 fixed addresses.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

When a lease expires, the DHCP server removes the A, AAAA, and PTR records that it updated. It does not remove any records that the client updated.

## Replacing Host Names for DDNS Updates

In situations where you need to restrict the use of specific characters in a host name for DDNS updates, you can configure a hostname rewrite policy. Such policy accepts certain characters and replaces others in host names specified in IPv4 DHCP requests. When you create a hostname rewrite policy, you enter a list of valid characters that the appliance accepts in the host name. You also specify a character that the appliance uses to replace invalid characters. You can create multiple hostname rewrite policies on the appliance, but you can only enable one policy at any given time. The appliance provides a default policy that includes **a-z0-9\_** as valid characters and dash (-) as the replacement character. You cannot modify or delete the default policy.

When you enable a hostname rewrite policy, the appliance replaces host names with the newly translated host name when it issues DHCP leases and sends DDNS updates for IPv4 DHCP clients. For information about how to add and enable a hostname rewrite policy, see [Adding and Enabling a Hostname Rewrite Policy](#) below.

Before you enable a hostname rewrite policy, consider the following:

- You must enable DDNS updates before the hostname rewrite policy can take effect.
- You can use a hostname rewrite policy only if MS code pages are disabled.
- The policy supports only IPv4 DHCP clients.
- If DHCP option 81 support is enabled and updating DDNS is in the request, the appliance sends updates for A records directly to the DNS server and DHCP only updates the PTR record. When this happens, there can be a mismatch in the host name between the A and PTR records.
- Changes made to a hostname rewrite policy apply only to subsequent DDNS updates.

When an IPv4 DHCP client requests an IP address, it includes its host name in DHCP option 12. If you enable a hostname rewrite policy, the appliance uses the newly translated host name when it issues a lease to the client. The client can also include a FQDN in option 81, in which it instructs the server whether to perform DDNS updates. If the client sends a FQDN in option 81, the appliance replaces the entire FQDN based on the policy. For example, if the FQDN in option 81 is dev.bldg12.corpxyz.com, the appliance replaces invalid characters in the entire FQDN even though the host name can be dev or dev.bldg12. For example, if your hostname rewrite policy specifies valid characters as **a-z** and the replacement character is -, the newly translated FQDN is dev.bldg--.corp---.com. For information about client FQDN in option 81, see [About the Client FQDN Option](#).

Note that when multiple IPv4 DHCP clients specify host names that map to the same translated host name, the appliance allocates leases to all clients, but it only sends DDNS updates to the first client request. When it tries to update DNS for subsequent clients, the updates fail.

You can add and enable a hostname rewrite policy for the entire Grid. You can also override the policy at a member level, as described in [Overriding a Grid Hostname Rewrite Policy](#) below.

## Adding and Enabling a Hostname Rewrite Policy

To add and enable a hostname rewrite policy, complete the following:

1. From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Properties**.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode**.
3. In the **IPv4 DDNS -> Advanced** tab, click the Add icon in the **Hostname Rewrite Policy** section:
  - **Policy Name:** Enter the policy name. Each policy name must be unique.
  - **Valid Characters:** Enter a list of valid characters you want to keep in the host name. Ensure that you consider the following rules:
    - You can include only printable ASCII characters and space.
    - The appliance includes period (.) as a valid character by default. You do not need to specify it.
    - You can also use shortcuts for a series or range of characters. For example, when you enter **a-d**, the appliance includes the following: A, B, C, D, a, b, c, and d. When you enter **0-5**, the appliance includes the following: 0, 1, 2, 3, 4, and 5. In a character range, ensure that the start character is less than the end character.
    - If you want to use dash (-) as a character, ensure that you put it in front of the valid character pattern. Otherwise, the appliance treats the string as a range of characters.
    - You can build a POSIX regular expression based on the string you enter here, but you cannot enter an empty string.
    - You cannot use the meta character (^) as a start or end character in a range. For example, **a-^** is invalid. You also cannot use duplicate characters as character sets. For example, **aa** is invalid.
  - **Replace Invalid Characters with:** Enter a character the appliance uses to replace invalid characters. Only enter one printable ASCII character. You cannot enter multiple characters or use space as the replacement character.

To test the hostname policy before adding it to the system, enter a sample hostname in the **Sample Host Name** field, and then click **Test**. The appliance displays the translated hostname. You can change the policy and test it again until you get the desired result. Click the Add icon to add the new hostname rewrite policy to the table. The appliance comes with a default policy that includes **a-z0-9\_** as valid characters and dash (-) as the replacement character. Grid Manager displays the following for each policy:

- **Policy Name:** The name of the hostname rewrite policy.
  - **Valid Characters:** Valid characters for the host name.
  - **Replace Invalid Characters with:** The character used to replace invalid characters in the host name. You can also select a hostname policy and click the Edit icon to modify it, or click the Delete icon to delete it. You cannot modify or delete the default policy. For information about how to modify a policy, see [Modifying a Hostname Rewrite Policy](#) below.
4. Complete the following to enable the hostname rewrite policy:
    - **Enable hostname rewrite policy:** Select this checkbox to use a hostname rewrite policy for DHCP leases and DDNS updates for IPv4 DHCP clients. From the drop-down list, select the hostname policy you want to use.
  5. Save the configuration.

## Modifying a Hostname Rewrite Policy

To modify a hostname rewrite policy, complete the following:

1. From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Properties**.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode**.
3. In the **IPv4 DDNS -> Advanced** tab, do the following in the **Hostname Rewrite Policy** section:
  - Select a policy from the table, and then click the Edit icon.
  - In the Edit Hostname Rewrite Policy section, modify and test the policy as described in [Adding and Enabling a Hostname Rewrite Policy](#) above.  
Note that if you enable the policy at the Grid level, you can modify all information, including the policy name. If you enable the policy at the member level, you can modify any information, except for the policy name.
4. Click **Save**. The appliance updates the policy in the table.
5. Save the configuration.

## Overriding a Grid Hostname Rewrite Policy

You can override a Grid hostname rewrite policy at the member level. To override a Grid policy, complete the following:

1. From the **Data Management** tab, select the **DHCP** tab and click the **Members** tab -> *member* checkbox -> *Edit* icon.
2. In the *Member DHCP Properties* editor, click **Toggle Advanced Mode**.
3. In the **IPv4 DDNS** -> **Advanced** tab, click **Override** in the **Hostname Rewrite Policy** section, and then complete the following:
  - **Enable hostname rewrite policy:** Select this checkbox to use a hostname rewrite policy, or deselect the checkbox to disable the policy.
  - When you enable this feature, select a policy that you want to use from the drop-down list. Grid Manager displays all hostname rewrite policies that you have configured on the appliance in the drop-down list. After you select a policy, Grid Manager displays the policy name, valid characters, and the replacement character.
4. Save the configuration.

## About the Client FQDN Option

When an IPv4 DHCP client sends DHCP DISCOVER and DHCP REQUEST messages, it can include option 81, the Client FQDN option. An IPv6 DHCP client can include option 39, the Client FQDN option, when it sends Solicit and Request messages.

The Client FQDN option contains the FQDN (fully qualified domain name) of the client and instructions on whether the client or the server performs DDNS updates. You can configure the appliance to replace the FQDN in the option by defining a hostname rewrite policy. For information about adding and enabling a hostname rewrite policy, see [Replacing Host Names for DDNS Updates](#).

The DHCP server can support option 81 for IPv4 and IPv6 clients, and use the host name or FQDN that the client provides for the update. It can also allow or deny the client's request to update DNS, according to the administrative policies of your organization. The DHCP server indicates its response in the DHCP OFFER message it sends back to an IPv4 client, and in the Reply message it sends back to an IPv6 client.

## Sending Updates with the FQDN Option Enabled

When you enable the DHCP server to support the FQDN option, it uses the information provided by the IPv4 or IPv6 client to update DNS as follows:

- When an IPv4 or IPv6 DHCP client sends a DHCP request with the FQDN option, it can include either its FQDN or only its host name.
  - If the request includes the FQDN, the DHCP server uses this FQDN to update DNS. You can specify a list of forward-mapping zones to be updated for IPv4 and IPv6 clients using the FQDN option. For information, see [Sending Updates for DHCP Clients Using the FQDN Option](#) below.
  - If the request includes the host name, the DHCP server provides the domain name. It combines the host name of the client and the domain name to create an FQDN for the client. It then updates DNS with the FQDN it created. (You can enter the domain name in the General page of the DHCP Properties window. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).)
- When a DHCP client sends a DHCP request with its hostname, the DHCP server adds the domain name you specified to create an FQDN for the client. It then updates DNS with the FQDN it created. For information about entering the domain name, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).
- When a DHCP client does not send a host name, the DHCP server provides a lease but does not update DNS. You can configure the DHCP server to generate a host name and update DNS as described in [Generating Host Names for DDNS Updates](#).
- If multiple DHCP clients specify the same FQDN or host name, the DHCP server allocates leases to the clients, but updates DNS only for the client that first sent the request. When it tries to update DNS for the succeeding clients, the update fails.

## Sending Updates from DHCP Clients or a DHCP Server

When you enable the DHCP server to support the FQDN option, you must decide if you want the DHCP server to allow clients to update DNS. If you allow the client to update DNS, then the client updates its A or AAAA record only. The DHCP server always updates the PTR records. You can configure the DHCP server as follows:

- The DHCP server can allow clients to update DNS when they send the request in the FQDN option. This is useful for small sites where security is not an issue or in sites where clients move from one administrative domain to another and want to maintain the same FQDN regardless of administrative domain.  
If you configure the DHCP server to allow clients to perform DDNS updates, you must also configure the DNS server to accept these updates from clients. Note that multiple clients can use the same name, resulting in multiple PTR records for one client name.  
When a lease expires, the DHCP server does not delete the A or AAAA record, if it was added by the client.
- The DHCP server can refuse the DHCP client's request to update DNS and always perform the updates itself. When the DHCP server updates DNS, it uses the FQDN provided by the DHCP client. Select this option if your organization requires tighter control over your network and does not allow clients to update their own records.

If you do not enable support for the FQDN option and a client includes it in a DHCP request with its FQDN, the DHCP server does not use the FQDN of the client. Instead, it creates the FQDN by combining the host name from the client with the domain name specified in the Grid or Member DHCP Configuration editor.

Do the following to configure support for the FQDN option for both IPv4 and IPv6 clients:

- Enable support for the option and specify who performs the DDNS update. For more information, see [Enabling FQDN Option Support](#) below.
- Specify the DNS zones and DNS view for the updates. For more information, see [Sending Updates for DHCP Clients Using the FQDN Option](#) below.

## Enabling FQDN Option Support

You can configure support for the FQDN option for IPv4 and IPv6 clients at the Grid, member, network and shared network levels.

To configure support for the FQDN Option (option 81) for IPv4 and (Option 39) for IPv6:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Properties**.  
**Member:** From the **Data Management** tab, select the **DHCP** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
**Network:** From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> *Networks* -> *network* checkbox -> Edit icon.  
**Shared Network:** From the **Data Management** tab, select the **DHCP** tab and click the **Networks** tab -> **Shared Networks** -> *shared\_network* checkbox -> Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the **IPv4 DDNS** -> **Advanced** tab for the Grid and member, or the **IPv4 DDNS** -> **Basic** tab for the network, do the following:
  - **Option 81 Support**
    - **Enable Option 81 Support:** Select this to enable the support for option 81.
    - **DHCP server always updates DNS:** Select this to allow the DHCP server to update DNS, regardless of the requests from DHCP clients.
    - **DHCP server updates DNS if requested by client:** Select this to allow the DHCP server to update DNS only when requested by DHCP clients.
3. In the **IPv6 DDNS** -> **Advanced** tab for the Grid and member, or the **IPv6 DDNS** -> **Basic** tab for the network, do the following:
  - **FQDN Support:** Select **Enable FQDN Support** and select one of the following to indicate whether the DHCP server or the client performs the DDNS update.
    - DHCP always updates DNS
    - DHCP updates DNS if requested by client
4. Save the configuration and click **Restart** if it appears at the top of the screen.

When a lease expires, the DHCP server removes the A or AAAA records and PTR records that it updated. It does not remove any records that the client updated.

## Sending Updates for DHCP Clients Using the FQDN Option

You must specify the DNS view to be updated for each network view.

To send updates to zones for DHCP IPv4 and IPv6 clients using the FQDN option:

1. If there are multiple network views in the Grid, select a network view.
2. From the **Data Management** tab, select the **DHCP** tab, and then click **Configure DDNS** from the Toolbar.
3. In the *DDNS Properties* editor, complete the following:
  - **DNS View:** If a network view has more than one DNS view, this field lists the associated DNS views. From the drop-down list, select the DNS view to which the DHCP server sends DDNS updates. Otherwise, the appliance uses the default DNS view.
  - **Zones to Update for Hosts Using DHCP FQDN Option:** In this section, you can define forward-mapping zones to which the DHCP server sends DDNS updates for IPv4 and IPv6 DHCP clients that use the FQDN option. You must first enable support for the FQDN option before the DHCP server can send DDNS updates to these zones. By default, the DHCP server sends DDNS updates to zones using the domain name that you define for DHCP objects, such as networks and DHCP ranges. For clients using this option, the DHCP server uses the domain name defined in the option.  
Click the Add icon to specify a forward-mapping zone. Note that the Forward-mapping Zone Selector dialog box displays only the DNS zones that are associated with the selected DNS view. The zones you select here are written to the `dhcpd.conf` file and the `dhcpdv6.conf` file as "zone" statements with the matching TSIG key of the DNS view, so the updates are sent to the correct DNS view.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring DDNS Update Verification

The DHCP server can handle DDNS updates differently, depending on how stringently you configure record handling. You can configure the DHCP server to update records only after passing verification. You can adjust the way DHCP handles updates so the DHCP server updates records after passing less stringent verification requirements, or without any type of verification.

To provide a measure of protection against unintentional changes of DNS data, NIOS appliances support the generation and use of TXT records, as described in IETF draft, *draft-ietf-dhc-dhcp-dns-12.txt* and by the ISC (Internet Systems Consortium). When DHCP updates or deletes an A or AAAA record, the corresponding TXT record is checked first to verify the authenticity of the update. The TXT record is based on a hash of the DHCID which is unique to each client, usually based in part on the MAC address or the DUID. If the client requests an update to DNS, the DHCP server first checks the TXT record to verify that it matches the client that originally inserted the record. This process provides assurance that the updates are from the same client. These security checks are based upon inserting a cryptographic hash of the DHCID (DHCP Client Identifier) into a DNS TXT RR and then verifying that value before updating. For example, a sample client update adds the following records in DNS:

oxcart.lo0.net.	21600	IN A 172.31.1.20
oxcart.lo0.net.	21600	IN TXT "313ce164780d34b91486b7c489ed7467e6"
20.1.31.172.in-addr.arpa.	21600	IN PTR oxcart.lo0.net.

However, your DNS configuration might require that the NIOS appliance handle DNS record updates differently than described in *draft-ietf-dhc-dhcp-dns-12.txt*. Your specific requirements might benefit from less-stringent verification of the DHCID, or might require skipping verification entirely. Verification checks might cause complications in some specific cases described below:

- **Mobility:** The TXT record is based on the DHCID unique to each client and is usually based on the MAC address or DUID of the interface. Devices such as laptops that connect to both wired and wireless networks have different MAC addresses or DUIDs and different DHCID values for each interface. In this scenario, after either one of the



network interfaces inserts a DNS record, updates are allowed from that interface only. This results in a disruption of service for DDNS updates when roaming between wired and wireless networks.

- **Migration:** The second problem occurs during a migration from non-ISC based systems to ISC systems. For example, if the user is migrating from a Microsoft-based system, the clients have A or AAAA and PTR records in the DDNS updates but no TXT records. As a result, new DDNS updates fail after the migration.
- **Mixed Environments:** The final problem occurs in mixed ISC and non-ISC environments. For example, assume that both Microsoft and ISC DHCP servers update DNS records on the appliance. In a mixed environment, since the Microsoft DHCP server does not insert the TXT records, DDNS updates from ISC-based systems fail while updates from the Microsoft DHCP server are committed into the database. This behavior is applicable only when you select **Standard ISC** and **Check TXT only** DDNS update verification modes.

The NIOS appliance offers four modes to handle DDNS updates as described in the [DDNS Update Verification Mode table](#) :

*DDNS Update Verification Mode*

Mode	If a Record at Lease Grant	Then TXT Record at Lease Grant	Lease Grant Action	Lease Expire Action
Standard ISC	Exists	Must match	Delete A or AAAA, TXT if exists Add A or AAAA Add PTR	Delete PTR Delete A or AAAA, TXT if TXT matches and no other A or AAAA RRs
	No A or AAAA record	No check	Add A or AAAA, TXT Add PTR	
Check TXT only	Exists	Must exist	Delete A or AAAA, TXT Add A or AAAA, TXT Add PTR	Delete PTR Delete A or AAAA if TXT exists and no other A or AAAA RRs
	No A or AAAA record	No check	Add A or AAAA, TXT Add PTR	
ISC Transitional	Exists	No check	Delete A or AAAA, TXT if exists Add A or AAAA, TXT Add PTR	Delete PTR Delete A or AAAA, TXT if TXT matches and no other A or AAAA RRs
	No A or AAAA record	No check	Add A or AAAA, TXT Add PTR	
No TXT record	Exists	No check	Delete A or AAAA Add A or AAAA Add PTR	Delete PTR, A or AAAA
	No A or AAAA record	No check	Add A or AAAA Add PTR	

Depending on your expected usage, you must carefully consider the various options for update verification. The following section illustrates recommendations for each verification option:

- **Standard ISC:** This method is the most stringent option for verification of updates. This is the default.
- **ISC Transitional:** This method is useful during migrations from systems that do not support the TXT record to systems that are ISC-based.
- **Check TXT only:** This method is useful for the roaming laptop scenario. The NIOS appliance checks that a TXT record exists, but does not check the value of the TXT record.
- **No TXT record:** This method should be used with caution because anyone can send DDNS updates and overwrite records. This method is useful when both ISC and non-ISC-based DHCP servers and clients are updating the same zone. Infoblox recommends that you allocate a DNS zone for this authentication method, as a precaution.



Note that in certain situations, when a DHCP lease expires, the DHCP server might remove the TXT record even if there is no A or AAAA record.

You can enable this feature at the Grid level. To configure TXT record handling on the DHCP server:

1. From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Properties**.
2. In the **IPv4 DDNS -> Advanced** tab or the **IPv6 DDNS -> Advanced** tab, select one of the following from the **TXT (DHCID) Record Handling** drop-down list:
  - **Check Only:** Select this checkbox to enable minimal checking of DDNS updates. Specifically, A or AAAA records are modified only if a TXT record exists. The NIOS appliance checks that a TXT record exists, but does not check its value.
  - **ISC:** Select this checkbox to enable standard ISC (Internet Systems Consortium) handling for DDNS updates. Specifically, A or AAAA records are modified or deleted only if the TXT records match. This option is the default setting on the appliance.
  - **ISC Transitional:** Select this checkbox to enable less stringent handling of DDNS updates. Specifically, the NIOS appliance enables you to add or modify A or AAAA records whether or not TXT records exist. It checks whether a TXT record exists and then processes the update. If the appliance does not find a TXT record, it adds the record.
  - **No TXT Record:** Select this checkbox to disable TXT record checking. Specifically, A or AAAA records are added, modified, or deleted whether or not the TXT records match. No TXT records are added, and existing TXT records are ignored.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring DNS Servers for DDNS

For security reasons, an Infoblox DNS server does not accept DDNS updates by default. You must specify the sources from which you want to allow the DNS server to receive updates. You can configure the Infoblox DNS server to receive updates from specified DHCP clients, as described in [Enabling DNS Servers to Accept DDNS Updates below](#), and to accept forwarded updates from another DNS server, as described in [Forwarding Updates below](#).

For protection against spoofed IP addresses, you can use TSIG (transaction signatures) to authenticate and verify updates.

TSIG uses the MD5 (Message Digest 5) algorithm and a shared secret key to create an HMAC (hashed message authentication code) — sometimes called a *digital fingerprint* — of each update. Both the DHCP server sending the update and the DNS server receiving it must share the same secret key. Also, it is important that the time stamps on the TSIG-authenticated updates and update responses be synchronized, or the participants reject them. Therefore, use an NTP server to set the time on all systems involved in TSIG authentication operations.

The TSIG key that you use can come from several places:

- You can use the key generation tool described in this section to create a new TSIG key to authenticate updates from the DHCP server.
- You can enter (copy and paste) a TSIG key that you previously generated for another purpose, such as for zone transfers.
- If the DHCP server is on a separate appliance and a TSIG key was previously generated on that appliance, you can enter (copy and paste) that TSIG key onto the local DNS server.

The TSIG key name and value that the DHCP and DNS servers use must be the same.



### Note

Whether you deploy NIOS appliances in a Grid or independently, they send updates to UDP port 53. Grid members do not send updates through a VPN tunnel. Grid members do, however, authenticate updates between them using TSIG (transaction signatures) based on an internal TSIG key.

## Enabling DNS Servers to Accept DDNS Updates

You can configure the Infoblox DNS server to receive updates from specified DHCP clients only. You can set this for the Grid so that the Grid members receive DDNS updates only from the specified sources. Note that you specify the IP

addresses of the sources of the updates and not the actual IP addresses in the DNS records being updated.  
To configure the DNS server to accept updates from the specified sources complete the following steps:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
**Zones:** From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab-> *dns\_view* -> *zone* checkbox -> Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the editor, click **Toggle Advanced Mode**, select the **Updates** tab.  
Ensure that you understand how the appliance handles match lists before you specify the list of IP sources for DDNS updates, as described in [You can use the following OpenStack cloud-init template to configure an IB-V815 as a Grid Master](#).
3. In the *Allow updates from* section, select one of the following:
  - **None:** Select this to deny DDNS updates from any DHCP clients. This is selected by default.
  - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance receives DDNS updates from the sources that have the **Allow** permission in the named ACL. You can click **Clear** to remove the selected named ACL.
  - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows:
    - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
    - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
      - **Permission:** Select **Allow** or **Deny** from the drop-down list.
    - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
      - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
      - **Permission:** Select **Allow** or **Deny** from the drop-down list.
    - **TSIG Key:** In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
      - **Key name:** Enter a meaningful name for the key, such as a zone name or the name of a remote name server. This name must match the name of the same TSIG key on other name servers.
      - **Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.
      - **Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop-down list, and then click **Generate Key Data** to create a new key. You must enable GSS-TSIG signed updates to receive DDNS updates from TSIG key based ACEs. For information about how to enable this, see [Accepting GSS-TSIG Updates](#).
  - **Any Address/Network:** Select this to receive DDNS updates from any IP addresses.  
After you have added access control entries, you can do the following:
    - Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
    - Reorder the list of ACEs using the up and down arrows next to the table.
    - Select an ACE and click the Edit icon to modify the entry.
    - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.

- **Allow GSS-TSIG signed updates:** This checkbox is selected only if you have enabled GSS-TSIG signed updates.
4. Optionally, you can:
    - Modify an item on the list by selecting it and clicking the Edit icon.
    - Remove an item from the list by selecting it and clicking the Delete icon.
    - Move an item up or down the list. Select it and drag it to its new position, or click the up or down arrow. The appliance applies permissions to items in the order they are listed.
  5. Save the configuration.

## Forwarding Updates

When a secondary DNS server receives DDNS updates, it must forward the updates to the primary server because it cannot update zone data itself. In such situations, you must enable the secondary server to receive updates from the DHCP server, and then forward them to the primary DNS server.

To configure the secondary server to accept and forward updates for all zones:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the **Toolbar** and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
**Zones:** From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns\_view* -> *zone* checkbox -> Edit icon.  
 To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. In the editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Advanced** subtab of the **Updates** tab, and then complete the following:
  - **Allow secondary name servers to forward updates:** Select this checkbox.
  - **Forward updates from:** This is available only for authoritative zones. Click **Add**. Depending on the item that you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows:
    - **None:** Select this to deny DDNS updates from any clients. This is selected by default.
    - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL. Grid Manager displays the *Named ACLs* Selector. Select the named ACL you want to use. If you have only one named ACL, Grid Manager automatically displays the named ACL. When you select this option, the appliance receives DDNS updates from the sources that have the **Allow** permission in the named ACL. You can click **Clear** to remove the selected named ACL.
    - **Set of ACEs:** Select this to configure individual ACEs. Click the Add icon and select one of the following from the drop-down list. Depending on the item you select, Grid Manager either adds a row for the selected item or expands the panel so you can specify additional information about the item you are adding, as follows.
      - **IPv4 Address and IPv6 Address:** Select this to add an IPv4 address or IPv6 address. Click the **Value** field and enter the IP address. The **Permission** column displays **Allow** by default. You can change it to **Deny** by clicking the field and selecting **Deny** from the drop-down list.
      - **IPv4 Network:** In the **Add IPv4 Network** panel, complete the following, and then click **Add** to add the network to the list:
        - **Address:** Enter an IPv4 network address and either type a netmask or move the slider to the desired netmask.
        - **Permission:** Select **Allow** or **Deny** from the drop-down list.
      - **IPv6 Network:** In the **Add IPv6 Network** panel, complete the following, and then click **Add** to add the network to the list:
        - **Address:** Enter an IPv6 network address and select the netmask from the drop-down list.
        - **Permission:** Select **Allow** or **Deny** from the drop-down list.
    - **TSIG Key:** In the **Add TSIG Key** panel, complete the following, and then click **Add** to add the TSIG key to the list:
      - **Key name:** Enter a meaningful name for the key, such as a zone name or the name of a remote name server. This name must match the name of the same TSIG key on other name servers.
      - **Key Algorithm:** Select either **HMAC-MD5** or **HMAC-SHA256**.
      - **Key Data:** To use an existing TSIG key, type or paste the key in the **Key Data** field. Alternatively, you can select the key algorithm, select the key length from the **Generate Key Data** drop-down list, and then click **Generate Key Data** to create a new key.

You must enable GSS-TSIG signed updates to receive DDNS updates from TSIG key based ACEs. For information about how to enable this, see [Accepting GSS-TSIG Updates](#).

- **Any Address/Network:** Select to allow or disallow the appliance to receive DDNS updates from any IP address.

After you have added access control entries, you can do the following:

- Select the ACEs that you want to consolidate and put into a new named ACL. Click the Create new named ACL icon and enter a name in the *Convert to NamedACL* dialog box. The appliance creates a new named ACL and adds it to the **Named ACL** panel. Note that the ACEs you configure for this operation stay intact.
- Reorder the list of ACEs using the up and down arrows next to the table.
- Select an ACE and click the Edit icon to modify the entry.
- Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.

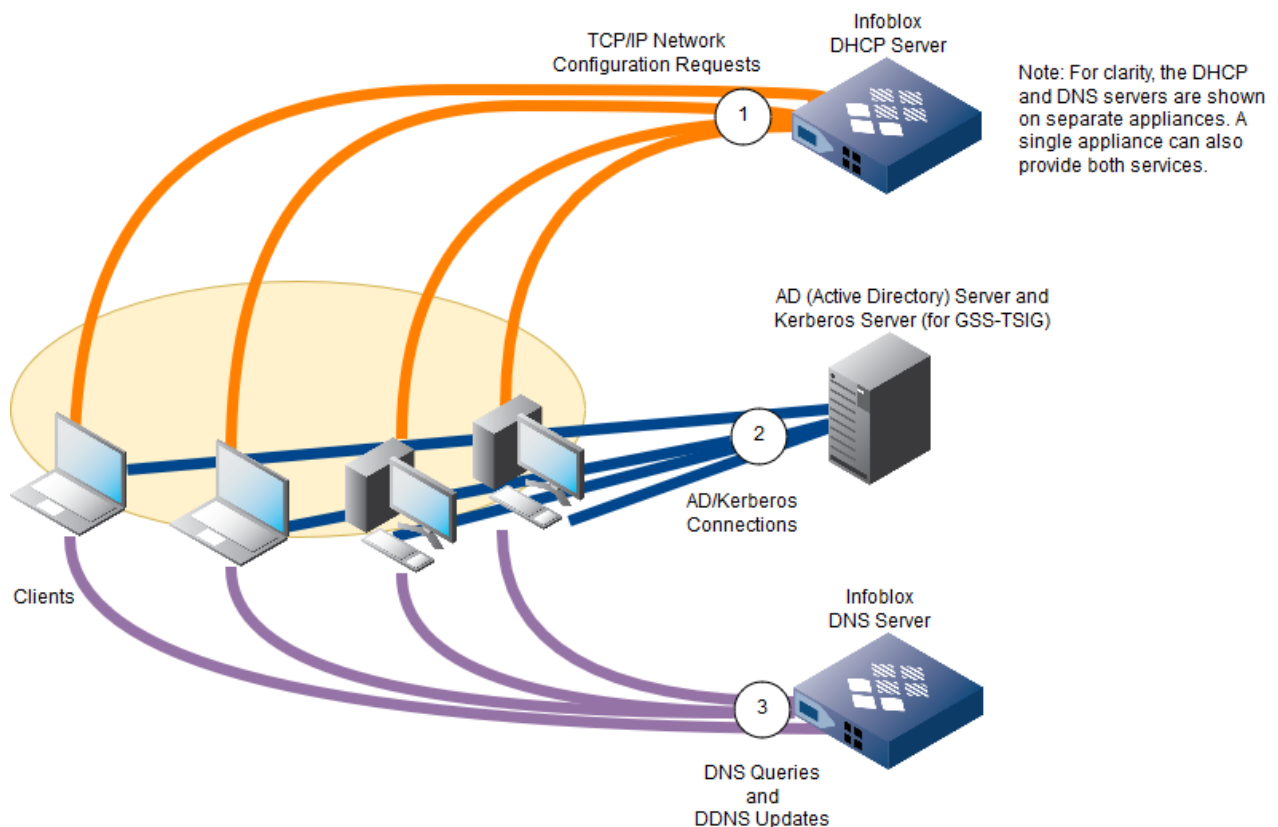
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Supporting Active Directory

Active Directory™ (AD) is a distributed directory service that authenticates network users and — by working with DHCP and DNS — provides the location of and authorizes access to services running on devices in a Windows® network. You can integrate a NIOS appliance providing DHCP and DNS services with servers running Windows 2000 Server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, or Windows Server 2016 with the Active Directory service installed. Assuming that you already have AD set up and it is currently in use, you can migrate DHCP and DNS services away from internal operations on the AD domain controller or from other third party DHCP and DNS systems to NIOS appliances that serve DHCP and DNS.

A NIOS appliance providing DHCP and DNS services to an AD environment can send and receive DDNS updates. In addition, a NIOS appliance can use GSS-TSIG (Generic Security Service-Transaction Signatures) authentication for DDNS updates. The basic DHCP, AD, and DNS services are shown in the [following DHCP, Active Directory, and DNS figure](#).

*DHCP, Active Directory, and DNS*



### Sending DDNS Updates to a DNS Server

You can configure an Infoblox DHCP server to send unauthenticated or GSS-TSIG-authenticated DDNS updates to a DNS server in an AD domain. There are no special configurations to consider when configuring a NIOS appliance to send unauthenticated DDNS updates to the DNS server. (For information about configuring DHCP, see [Configuring DHCP Properties](#), and for information on configuring the DHCP server to send DDNS updates, see [Configuring DHCP for DDNS](#).) For information about configuring a DHCP server to send GSS-TSIG authenticated updates, see [About GSS-TSIG](#).

### About GSS-TSIG

GSS-TSIG (Generic Security Service Algorithm for Secret Key Transaction) is used to authenticate DDNS updates. It is a modified form of TSIG authentication that uses the Kerberos v5 authentication system. GSS-TSIG involves a set of client/server negotiations to establish a "security context." It makes use of a Kerberos server (running on the AD domain controller) that functions as the KDC (Kerberos Key Distribution Center) and provides session tickets and temporary session keys to users and computers within an Active Directory domain. The client and server collaboratively create and mutually verify transaction signatures on messages that they exchange. Windows 2000 server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016 and Windows Server 2019 all support DDNS updates using GSS-TSIG. You can configure the appliance to accept GSS-TSIG signed DDNS updates from a single client or multiple clients that belong to different AD domains in which each domain has a unique GSS-TSIG key. You can also configure the appliance to support one or multiple GSS-TSIG keys for each Grid member. For information about how to configure GSS-TSIG for DHCP and DNS, see [Configuring GSS-TSIG keys](#). This feature also supports HA pairs and is compatible with DNS zones that have multiple primary servers configured. For more information about HA pairs and DNS zones with multiple primary servers, see [About HA Pairs](#) and [Assigning Zone Authority to Name Servers](#) respectively. You can upload keytab files that contain one or multiple GSS-TSIG keys and manage the keys globally. NIOS supports up to 256 GSS-TSIG keys for each member in the Grid. NIOS logs administrative changes to GSS-TSIG keys in the audit

log and failures in parsing or loading the keytab files in the syslog. Note that this feature is enabled only when you have installed the DNS license.



#### Note

For information about GSS-TSIG, see *RFC 3645, Generic Security Service Algorithm for Secret Key Transaction Authentication for DNS (GSS-TSIG)*.

A NIOS appliance can use GSS-TSIG authentication for DDNS updates for either one of the following:

- A NIOS appliance serving DHCP can send GSS-TSIG authenticated DDNS updates to a DNS server in an AD domain or multiple AD domains whose domain controller is running Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016 or Windows Server 2019. The DNS server can be in the same AD domain as the DHCP server or in a different domain.
  - For information about sending secure DDNS updates to a DNS server in the same domain, see [Sending Secure DDNS Updates to a DNS Server in the Same Domain](#) below.
  - For information about sending secure DDNS updates to a DNS server in a different domain, see [Sending Secure DDNS Updates to a DNS Server in Another Domain](#) below.
- A NIOS appliance serving DNS can accept GSS-TSIG authenticated DDNS updates from DHCP clients and servers in an AD domain or multiple AD domains whose domain controller is running Windows 2000 Server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016 or Windows Server 2019.
  - For information, see [Accepting GSS-TSIG-Authenticated Updates](#).  
Note that a NIOS appliance cannot support both of these features at the same time.

## Kerberos Authentication for GSS-TSIG

A keytab file contains pairs of Kerberos principal names and their corresponding encryption keys. It can contain keys for a single realm or multiple realms. It is possible to infer the KDC from the principal because Windows uses uppercase AD domain names for Kerberos realm names. You must provide the principal name. The principal name may contain Kerberos realm, and the DNS servers for the domain are available for DNS name resolution. Therefore, resolving `SRV_kerberos._tcp.REALM` will return the appropriate KDC. New TGTs cannot be acquired when the KDC that issues the TGT fails. If the appliance has successfully authenticated before the KDC failure, the secure updates will continue until the session key and TGT expire. The default expiration on Windows is 10 hours. If the appliance restarts or reboots, secure updates are deferred until the KDC becomes available.

Infoblox recommends restarting the DHCP service on NIOS to avoid any update failures, if the encryption key type is changed on the Microsoft server.

The following provides information about the traffic flow between the appliance and the KDC:

- Client uses keytab to get TGT for principal from KDC (AS-REQ/AS-REP).
- Client uses TGT to get session ticket from KDC (TGS-REQ/TGS-REP).
- Client uses session ticket to acquire TKEY from DNS server (TKEY/TKEY).
- Client uses TKEY to sign DNS updates (DNS-TSIG/DNS-TSIG).

The DNS server authenticates into the domain when the keytab file is generated on the KDC and its SPN (Service Principal Name) is mapped to an account. The server's private key is known to itself and to the KDC. The KDC generates the ticket and the DNS server allows the update.

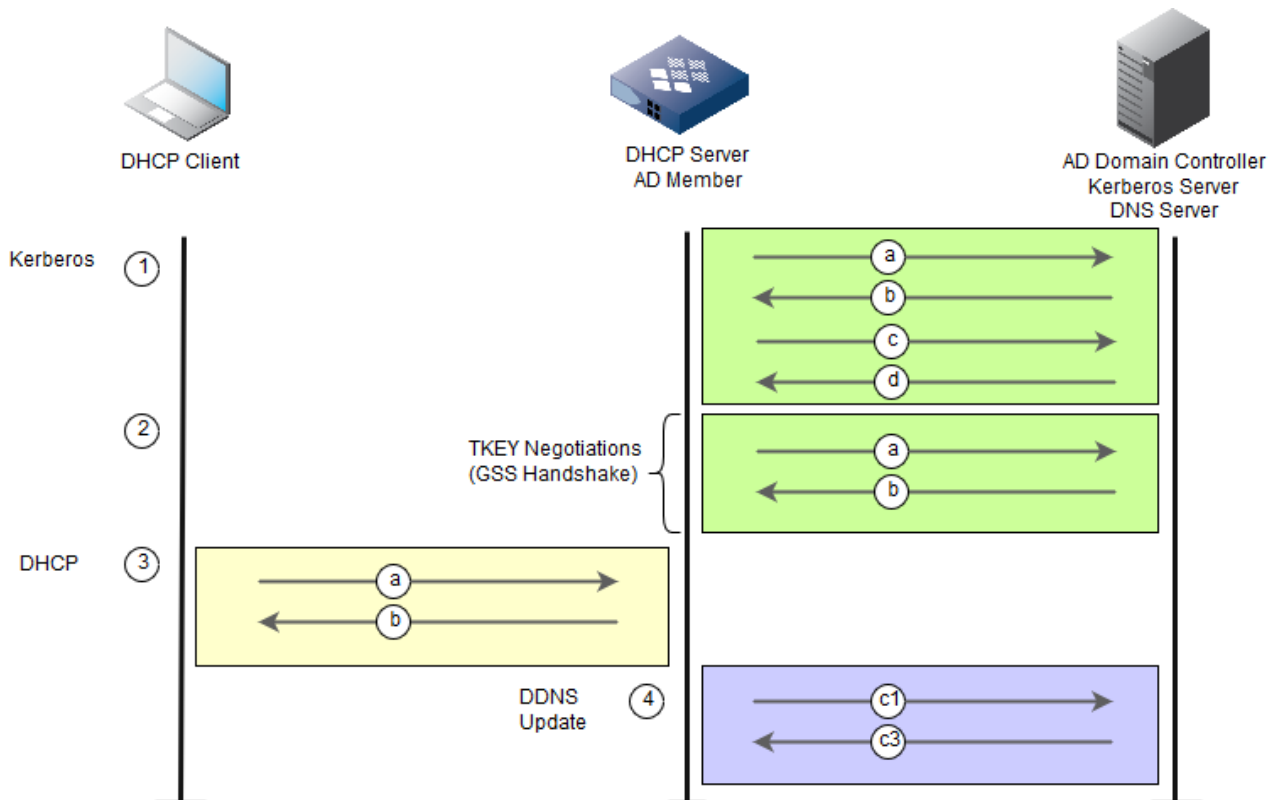
Note the following when you upload multiple keytab files on the appliance:

- NIOS displays an error message and discards the keytab file if the file does not have a recognizable key, SPN, version or encryption type, and it saves the error message in the syslog.
- NIOS considers duplicate keys as invalid keys if the keys have the same SPN, version, and encryption type.
- If NIOS encounters an invalid key during an upload, it will not upload the other keys in the keytab and the operation fails. NIOS saves the warning and error message in the syslog and in Grid Manager.

## Sending Secure DDNS Updates to a DNS Server in the Same Domain

An Infoblox DHCP server can send GSS-TSIG authenticated DDNS updates to a DNS server in an AD domain whose domain controller is running Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016 or Windows Server 2019. The DHCP server, DNS server, and domain controller are all in the same AD domain. The process by which an Infoblox DHCP server dynamically updates resource records on a DNS server using GSS-TSIG authentication is shown in the below figure. In the illustration, the Kerberos Key Distribution Center (KDC) is running on an AD domain controller, which also provides DNS service.

*An Infoblox DHCP Server Sends GSS-TSIG Updates to a DNS Server*



After you enable the NIOS appliance to send GSS-TSIG authenticated updates to a DNS server, the following process occurs:

1. Kerberos – Login, and TGT and Service Ticket Assignments
  - a. The Infoblox appliance automatically logs in to the AD/Kerberos server.
  - b. The Kerberos server sends the appliance a TGT (ticket-granting ticket).
  - c. Using the TGT, the appliance requests a service ticket for the DNS server.
  - d. The Kerberos server replies with a service ticket for that server.
2. TKEY negotiations (GSS Handshake):
  - a. The appliance sends the DNS server a TKEY (transaction key) request. A Transaction Key record establishes shared secret keys for use with the TSIG resource record. For more information, see *RFC 2930, Secret Key Establishment for DNS (TKEY RR)*. The request includes the service ticket. The service ticket includes the appliance's principal and proposed TSIG (transaction signature) key, along with other items such as a ticket lifetime and a timestamp.
  - b. The DNS server responds with a DNS server-signed TSIG, which is a "meta-record" that is never cached and never appears in zone data. A TSIG record is a signature of the update using an HMAC-MD5 hash that provides transaction-level authentication. For more information, see *RFC 2845, Secret Key Transaction Authentication for DNS (TSIG)*.

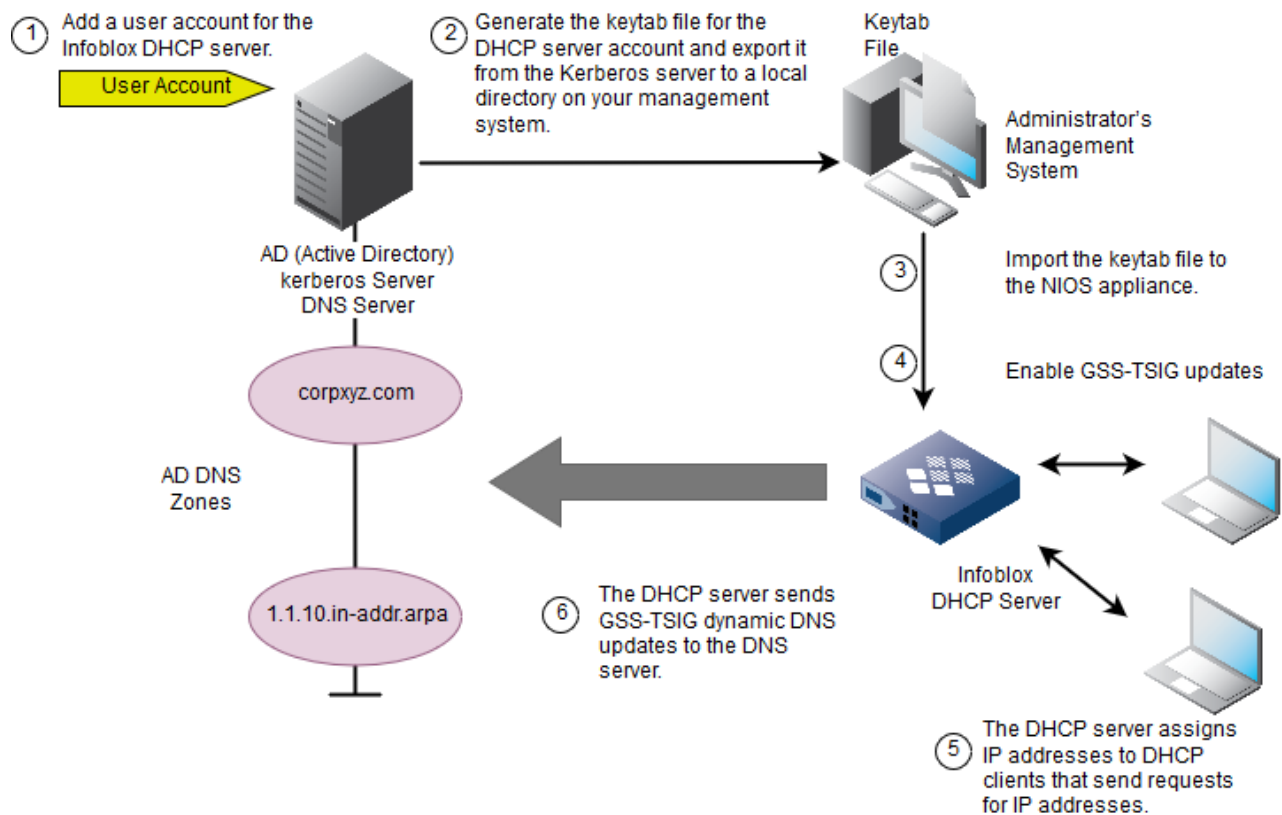


- The two participants have established a security context.  
When a DHCP client sends a request for an IP address to the DHCP server, the following occurs:
3. DHCP – IP Address and Network Parameters Assignment
    - a. The DHCP client requests an IP address.
    - b. The DHCP server assigns an IP address, subnet mask, gateway address, DNS server address, and a domain name.  
After the appliance assigns an IP address to the DHCP client, it sends the DDNS update to the DNS server as follows:
  4. DDNS – Dynamic Update of the Client's Resource Records
    - a. GSS-TSIG-Authenticated DDNS Update
      - i. The appliance sends an authenticated DDNS update, which may include the following resource records:
        - A or AAAA – Address record  
or  
PTR – Pointer record
        - TKEY – Transaction Key record
        - TSIG – TSIG record
      - ii. The DNS server verifies the DDNS update and allows it to complete.
      - iii. The DNS server sends a GSS-TSIG-authenticated response to the appliance, confirming the update.

### Configuring DHCP to Send GSS-TSIG Updates in the Same Domain

Before configuring an Infoblox DHCP server to support GSS-TSIG, you must create a user account on the Kerberos server for the appliance. Then you must export the corresponding keytab file from the Kerberos server and import it onto the NIOS appliance. The below figure illustrates the initial configuration tasks.

#### Adding an Infoblox DHCP Server to an AD Environment with GSS-TSIG Support





The Infoblox DHCP server can send GSS-TSIG-signed DDNS updates to a DNS server for one domain only, though multiple Infoblox DHCP servers can update that domain. If you want more than one Infoblox DHCP server to update a DNS domain, you can either import the same keytab file to the other Infoblox DHCP servers or generate and import a different keytab file. In a Grid, each member can update a different domain.



#### Note

For GSS-TSIG authentication to work properly, the system clock times of the Infoblox DHCP server, AD domain controller and DNS server must be synchronized. One approach is to use NTP and synchronize all three devices with the same NTP servers.

To use an AD domain controller as a Kerberos Key Distribution Center, complete the following tasks on an AD/Kerberos server:

1. Add a user account for the NIOS appliance to the AD domain controller. For information, see [Creating an AD User Account](#) below.
2. Generate the keytab file for the NIOS appliance account and export it from the AD domain controller to a local directory on your management system. For information, see [Generating and Exporting the Keytab File](#) below.

To configure a NIOS appliance to support AD and send GSS-TSIG secure DDNS updates to a DNS server, complete the following tasks on a NIOS appliance:

1. Import the keytab file from your management system to the appliance and enable GSS-TSIG dynamic updates at the Grid or member level. For information, see [Enabling GSS-TSIG Authentication for DHCP](#).
2. Configure the appliance to send GSS-TSIG dynamic updates to forward-mapping and optionally, reverse-mapping zones on the DNS server. For information, see [Managing GSS-TSIG keys](#).

## Creating an AD User Account

Connect to the AD domain controller and create a user account for the NIOS appliance.



#### Note

The name that you enter in the User logon name is the name that you later use when exporting the keytab file. This is also the principal name. The text in the First name, Initials, Last name, and Full name fields is irrelevant to this task.

The AD domain controller automatically creates a Kerberos account for this user. Note the following:

- If you define an expiration date for the user account and you later create a new account when the first one expires, the keytab for the corresponding Kerberos account changes. At that point, you must update the keytab file on the NIOS appliance (see [Generating and Exporting the Keytab File](#) below and [Enabling GSS-TSIG Authentication for DHCP](#)). Optionally, if your security policy allows it, you can set the user account for the NIOS appliance so that it never expires.
- If the AD domain controller is running Windows Server 2003, the user account must have the DES encryption type enabled. You can enable this either in the Account tab of the AD domain controller when you create the user account or by specifying **+DesOnly** when you use the Ktpass tool to generate the keytab file. For instructions, see the next section, [Generating and Exporting the Keytab File](#) below.
- The newly created AD user account must be a member of the DnsUpdateProxy group or an account that allows it to update records that have potentially been added by another DHCP server, such as DNS Admins.

## Generating and Exporting the Keytab File

You can use the Ktpass tool to generate and export the keytab file for the Kerberos account. Note that the version of the Ktpass tool that you use must match the Windows version of the domain controller. For example, if you are using a domain controller running Windows Server 2008 or Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016 or Windows Server 2019, you must use the Ktpass tool for Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016 or Windows Server

2019.

You enter different commands for generating and exporting the keytab file, depending on whether you are generating the keytab file from a server running Microsoft Windows 2000, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2010, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016 or Windows Server 2019.

A Windows Server 2003 domain controller allows you to generate a keytab file with only one key for a principal. A Windows Server 2008, Windows Server 2008 R2 domain controller, Windows Server 2012, Windows Server 2012 R2, Windows Server 2016 or Windows Server 2019 allows you to generate a keytab file with multiple keys for one principal. This is useful when the KDC has principals with multiple encryption types. When the NIOS DHCP server uses a keytab with multiple keys, it negotiates a key based on those in the configured keytab file.



**Note**

The keytab file contains highly sensitive data for the NIOS appliance account. Ensure that you store and transport its contents securely.

Infoblox strongly recommends the following encryption types for compatibility purposes:

Microsoft Windows Server	Export keytab
Microsoft Windows 2000	Specify <code>/crypto DES-CBC-MD5</code> as the export keytab.
Microsoft Windows 2003	Specify <code>/crypto RC4-HMAC-NT</code> as the export keytab. <ul style="list-style-type: none"><li>• Infoblox recommends that you do not use DES, but it is supported if you need it for compatibility with non-Windows systems.</li></ul>
Microsoft Windows 2008 and higher	Specify <code>/crypto RC4-HMAC-NT</code> as the export keytab. <ul style="list-style-type: none"><li>• You can also use AES, but RC4 is set by default for Windows 2008 servers.</li><li>• Infoblox recommends that you do not use DES, but it is supported if you need it for compatibility with non-Windows systems.</li></ul>

### Generating the Keytab on Windows 2000 Servers

To export the keytab file using a Microsoft Windows 2000 Resource Kit:

1. Start a command prompt.

2. Enter the following command to export the keytab file **for** the NIOS appliance user account:

```
C:> ktpass -princ service_name/FQDN_instance@REALM -mapuser AD_username -pass password -out filename.keytab
```

Note that the values are case-sensitive.

where:

- *service\_name/instance*: The AD user name for the NIOS appliance and a character string. The AD user name must match the user logon name on the AD domain controller.
- *REALM*: The Kerberos realm in uppercase. It must match the realm (or domain name) specified in the `-mapuser` option.

For example:

```
C:> ktpass -princ DNS/ns1.corpxyz.com@corpxyz.COM -mapuser ns1@corpxyz.com -pass
37Le37
-out ns1.keytab
```

### Generating the Keytab on Windows Servers 2003

The Ktpass tool is included in the Windows Server 2003 Support Tools. To export the keytab file using a Microsoft Windows 2003 Resource Kit:

1. Start a command prompt.
2. Enter the following command to generate the keytab file for the NIOS appliance user account:  
**ktpass -princ *service\_name/FQDN\_instance@REALM* -mapuser *AD\_username@REALM* -pass *password* -out *filename.ktb* -ptype KRB5\_NT\_PRINCIPAL -crypto RC4-HMAC-NT**  
Note that the values are case-sensitive.

The following are some examples of keytab file:

```
ktpass -princ HOST/ns1.corpxyz.com@GSS.LOCAL -mapuser gssuser@GSS.LOCAL -pass 37Le37
-out ns1.keytab -ptype krb5_nt_principal -crypto all
ktpass -princ gssuser@GSS.LOCAL -mapuser gssuser@GSS.LOCAL -pass 37Le37 -out
gssuser.keytab -ptype krb5_nt_principal -crypto all
ktpass -princ DNS/ns1.corpxyz.com@GSS.LOCAL -mapuser jsmith@GSS.LOCAL -pass 37Le37
-out ns1.ktb -ptype KRB5_NT_PRINCIPAL -crypto des-cbc-md5 +DesOnly
```

where

**-princ** = Kerberos principal. Note that this parameter is case sensitive. Specifies the principal name for the host or service in this format: DNS/ns1.corpxyz.com@GSS.LOCAL

- DNS = Service name in uppercase format.
- ns1.corpxyz.com = Instance in FQDN (fully-qualified domain name) format; this is the same as the DNS name of the NIOS appliance.
- GSS.LOCAL = The Kerberos realm in uppercase format. This must be the same as the AD domain name.

**-mapuser** = Maps the Kerberos principal name to the AD user account. If you omit the account name, mapping is deleted from the specified principal. You can use ksetup without any parameters or arguments to see the current mapped settings and the default realm. Example: ksetup /mapuser <Principal> <Account>. To create an AD user account, see Creating an AD User Account above.

- jsmith = The AD user name for the NIOS appliance.
- GSS.LOCAL = The Kerberos realm in uppercase. The realm (or domain name) must be the same as that specified in the **-princ** option.

**-pass** = The AD user account password. The Ktpass command changes the account password to the specified value, thus incrementing the version number of the user account and the resulting keytab file.

- 37Le37 = The password of the user account for the NIOS appliance.

**-out** = The name of the keytab file that is generated.


- ns1.ktb = The name of the keytab file

**-ptype** = Sets the principal type. This must be **krb5\_nt\_principal**.

**-crypto** = Specifies the encryption type. You can use the following encryption types:

- DES-CBC-CRC = Specifies DES encryption for the account. This encryption type is used for compatibility.
- DES-CBC-MD5 = Specifies DES encryption for the account. This encryption type adheres to the MIT implementation and is used for compatibility.
- RC4-HMAC-NT = Specifies 128-bit RC4-HMAC encryption for the account. This is enabled by default.

**+DesOnly** = Specifies DES encryption for the account.

 **Note**  
Windows Server 2003 does not support AES encryption.

After you execute the command to generate the keytab file, the AD domain controller displays a series of messages similar to the following to confirm that it successfully generated the keytab file:

```
Targeting domain controller: ibtest-xu5nxd56.corpxyz.local

Using legacy password setting method

Successfully mapped dns/anywhere to dns.

Key created.

Output keytab to dns.ktb:

Keytab version: 0x502

keysize 56 dns/anywhere@GSS.LOCAL ptype 1 (KRB5_NT_PRINCIPAL) vno 5 etype 0x3 (DES-
CBC-MD5)

keylength 8 (0xbae610f11552c80b)
```

Generating the Keytab on Windows Servers 2008 or Windows Servers 2008 R2

To generate the keytab file using the Ktpass tool:

1. Start a command prompt.
2. Enter the following command to generate the keytab file **for** the NIOS appliance user account:  
ktpass -princ username@REALM -mapuser logon\_name@REALM -pass password -out my.tab -ptype krb5\_nt\_principal -crypto encryption

Example:

```
ktpass -princ DNS/ns1.corpxyz.com@GSS.LOCAL -mapuser jsmith@GSS.LOCAL -pass 37Le37
-out ns1.keytab -ptype krb5_nt_principal -crypto RC4-HMAC-NT
```

where:

**-princ** = Kerberos principal. Note that this parameter is case sensitive. Specifies the principal name for the host or service in this format: DNS/ns1.corpxyz.com@GSS.LOCAL

- DNS = Service name in uppercase format.
- ns1.corpxyz.com = Instance in FQDN (fully-qualified domain name) format; this is the same as the DNS name of the NIOS appliance.
- GSS.LOCAL = The Kerberos realm in uppercase format. This must be the same as the AD domain name.

**-mapuser** = Maps the Kerberos principal name to the AD user account. If you omit the account name, mapping is deleted from the specified principal. You can use ksetup without any parameters or arguments to see the current mapped

settings and the default realm. Example: `ksetup /mapuser <Principal> <Account>`. To create an AD user account, see [Creating an AD User Account](#) above.

- `jsmith` = The AD user name for the NIOS appliance.
- `GSS.LOCAL` = The Kerberos realm in uppercase. The realm (or domain name) must be the same as that specified in the `-princ` option.

`-pass` = The AD user account password. The `Ktpass` command changes the account password to the specified value, thus incrementing the version number of the user account and the resulting keytab file.

- `37Le37` = The password of the user account for the NIOS appliance.

`-out` = The name of the keytab file that is generated.

- `ns1.ktb` = The name of the keytab file

`-ptype` = Sets the principal type. This must be `krb5_nt_principal`.

`-crypto` = Specifies the encryption type. Note that the **RC4-HMAC-NT** encryption type is enabled by default. You can also use the following:

- **DES-CBC-CRC** = Specifies DES encryption for the account. This encryption type is used for compatibility.
- **DES-CBC-MD5** = Specifies DES encryption for the account. This encryption type adheres to the MIT implementation and is used for compatibility.
- **RC4-HMAC-NT** = Specifies 128-bit RC4-HMAC encryption for the account. This is enabled by default.
- **AES256-SHA1** = Specifies 256-bit AES encryption for the account.
- **AES128-SHA1** = Specifies 128-bit AES encryption for the account.
- **ALL** = Specifies all of the above encryption types. Do not use this option if DES support is disabled.

You can optionally specify the following:

**+DesOnly** = Specifies DES encryption for the account. You must use this only when you use DES-CBC-MD5 for compatibility. Note that Windows 7 and Windows Server 2008 R2 do not support DES by default. However, you can enable DES on the Windows 2008 server. Include this option if you did not enable DES encryption for the account. For more information, refer to the information available in a third-party portal at: <http://weblogic-wonders.com/weblogic/2010/11/30/windows-7-des-encryption-support-for-kerberos-authentication/>



Note

You must not use `+Desonly` with `/crypto` all or other non-DES encryption types.

- **+setpass** = Sets a new AD user account password. This is required if the **+DesOnly** option is specified. When you use this encryption type, you must change the user's password. Otherwise, the ticket issued for the principal becomes unusable.

After you execute the command to generate the keytab file, the AD domain controller displays a series of messages similar to the following to confirm that it successfully generated the keytab file:

```
Targeting domain controller: qacert.test.local

Using legacy password setting method

Successfully mapped DNS/ns1.corpxyz.com to ns1.

Key created.

Output keytab to ns1.ktb: Keytab version: 0x502

keysize 80 DNS/ns1.corpxyz.com@GSS.LOCAL ptype 1 (KRB5_NT_PRINCIPAL) vno 3 etype 0x12
(AES256-SHA1)

keylength 32 (0xea8675d7abf13fd760a744088642fb917ceb6c9d267f5c54e595597846f06407)
```

## Generating the Keytab on Windows Servers 2012 or Windows Server 2012 R2

To generate the keytab file using the Ktpass tool:

- Start a command prompt.

- Enter the following command to generate the keytab file **for** the NIOS appliance user account:

```
ktpass -princ username@REALM -mapuser logon_name@REALM -pass password -out
my.tab -ptype krb5_nt_principal -crypto encryption
```

Example:

```
ktpass -princ DNS/ns1.corpxyz.com@GSS.LOCAL -mapuser jsmith@GSS.LOCAL -pass
37Le37 -out ns1.keytab -ptype krb5_nt_principal -crypto RC4-HMAC-NT
```

where:

**-princ** = Kerberos principal. Note that this parameter is case sensitive. Specifies the principal name for the host or service in this format: DNS/ns1.corpxyz.com@GSS.LOCAL

- DNS = Service name in uppercase format.
- ns1.corpxyz.com = Instance in FQDN (fully-qualified domain name) format; this is the same as the DNS name of the NIOS appliance.
- GSS.LOCAL = The Kerberos realm in uppercase format. This must be the same as the AD domain name.

**-mapuser** = Maps the Kerberos principal name to the AD user account. If you omit the account name, mapping is deleted from the specified principal. You can use ksetup without any parameters or arguments to see the current mapped settings and the default realm. Example: ksetup /mapuser <Principal> <Account>. To create an AD user account, see [Creating an AD User Account](#) above.

- jsmith = The AD user name for the NIOS appliance.
- GSS.LOCAL = The Kerberos realm in uppercase. The realm (or domain name) must be the same as that specified in the **-princ** option.

**-pass** = The AD user account password. The Ktpass command changes the account password to the specified value, thus incrementing the version number of the user account and the resulting keytab file.

- 37Le37 = The password of the user account for the NIOS appliance.

**-out** = The name of the keytab file that is generated.

- ns1.ktb = The name of the keytab file

**-ptype** = Sets the principal type. This must be **krb5\_nt\_principal**.

**-crypto** = Specifies the encryption type. You can specify the following encryption types:

- **DES-CBC-CRC** = Specifies DES encryption for the account. This encryption type is used for compatibility.
- **DES-CBC-MD5** = Specifies DES encryption for the account. This encryption type adheres to the MIT implementation and is used for compatibility.
- **RC4-HMAC-NT** = Specifies 128-bit RC4-HMAC encryption for the account. This is enabled by default.
- **AES256-SHA1** = Specifies 256-bit AES encryption for the account.
- **AES128-SHA1** = Specifies 128-bit AES encryption for the account.
- **ALL** = Specifies all of the above encryption types. Do not use this option if DES support is disabled.

After you execute the command to generate the keytab file, the AD domain controller displays a series of messages similar to the following to confirm that it successfully generated the keytab file:

```
Targeting domain controller: qacert.test.local

Using legacy password setting method

Successfully mapped DNS/ns1.corpxyz.com to ns1.

Key created.

Output keytab to ns1.keytab: Keytab version: 0x502

keysize 80 DNS/ns1.corpxyz.com@GSS.LOCAL ptype 1 (KRB5_NT_PRINCIPAL) vno 3 etype 0x12
(AES256-SHA1)

keylength 32 (0xea8675d7abf13fd760a744088642fb917ceb6c9d267f5c54e595597846f06407)
```

Generating the Keytab on Windows Servers 2016 or Windows Servers 2019

To generate the keytab file using the Ktpass tool:

1. Start a command prompt.
2. Enter the following command to generate the keytab file for the NIOS appliance user account:

```
ktpass -princ username@REALM -mapuser logon_name@REALM -pass password -out my.tab -ptype  
krb5_nt_principal -crypto encryption
```

Example:

```
ktpass -princ DNS/ns1.corpxyz.com@GSS.LOCAL -mapuser jsmith@GSS.LOCAL -pass  
37Le37 -out ns1.keytab -ptype krb5_nt_principal -crypto RC4-HMAC-NT
```

where:

**-princ** = Kerberos principal. Note that this parameter is case sensitive. Specifies the principal name for the host or service in this format: `DNS/ns1.corpxyz.com@GSS.LOCAL`

- **DNS** = This is an example of the service name in uppercase format.
- **ns1.corpxyz.com** = This is an example of the instance in FQDN (fully-qualified domain name) format; this is the same as the DNS name of the NIOS appliance.
- **GSS.LOCAL** = This is an example of the Kerberos realm in uppercase format. This must be the same as the AD domain name.

**-mapuser** = Maps the Kerberos principal name to the AD user account. If you omit the account name, mapping is deleted from the specified principal. You can use `ksetup` without any parameters or arguments to see the current settings and the default realm. Example: `ksetup /mapuser <Principal> <Account>`. To create an AD user account, see [Creating an AD User Account](#) above.

- `jsmith` = This is an example of the AD user name for the NIOS appliance.
- `GSS.LOCAL` = This is an example of the Kerberos realm in uppercase. The realm (or domain name) must be the same as that specified in the `-princ` option.

`-pass` = The AD user account password. The `Ktpass` command changes the account password to the specified value, thus incrementing the version number of the user account and the resulting keytab file.

- `37Le37` = This is an example of the password of the user account for the NIOS appliance.

`-out` = The name of the keytab file that is generated.

- `ns1.ktb` = This is an example of the name of the keytab file.

`-ptype` = Sets the principal type. This must be `krb5_nt_principal`.

`-crypto` = Specifies the encryption type. You can specify the following encryption types:

- **DES-CBC-CRC** = Specifies DES encryption for the account. This encryption type is used for compatibility purposes.
- **DES-CBC-MD5** = Specifies DES encryption for the account. This encryption type adheres to the MIT implementation and is used for compatibility purposes.
- **RC4-HMAC-NT** = Specifies 128-bit RC4-HMAC encryption for the account. This is enabled by default.
- **AES256-SHA1** = Specifies 256-bit AES encryption for the account.
- **AES128-SHA1** = Specifies 128-bit AES encryption for the account.
- **ALL** = Specifies all of the above encryption types. Do not use this option if DES support is disabled.

After you execute the command to generate the keytab file, the AD domain controller displays a series of messages similar to the following to confirm that it successfully generated the keytab file:

```
Targeting domain controller: qacert.test.local

Using legacy password setting method

Successfully mapped DNS/ns1.corpxyz.com to ns1.

Key created.

Output keytab to ns1.keytab:

Keytab version: 0x502

keysize 80 DNS/ns1.corpxyz.com@GSS.LOCAL ptype 1 (KRB5_NT_PRINCIPAL) vno 3 etype 0x12
(AES256-SHA1)

keylength 32 (0xea8675d7abf13fd760a744088642fb917ceb6c9d267f5c54e595597846f06407)
```

### Creating an External Zone for GSS-TSIG Updates

For each network view, specify the zone to be updated, the IP address of the primary DNS server for that zone, and the security method, GSS-TSIG. The zone must be in the same Active Directory domain as the member that is sending the updates.

You can add information for a forward and reverse zone. The DHCP server updates the A record in the forward zone and the PTR record in the reverse zone.

To enable the NIOS appliance to send dynamic updates to a DNS server using GSS-TSIG for authentication:

1. If there are multiple network views in the Grid, select a network view.
2. On the **Data Management** tab, select the **DHCP** tab, expand the Toolbar, and then click **Configure DDNS**.
3. In the *DNS Updates to External Zones* table of the *DDNS Properties* editor, click the Add icon and complete the following fields in the *Add External DDNS Zone* panel:



- **Zone Name:** Enter the name of the zone that receives the updates. You can specify both forward-mapping and reverse-mapping zones.
- **DNS Server Address:** Enter the IP address of the primary name server for that zone.
- **Security:** Choose **GSS-TSIG** from the drop-down list, complete the following, and then click **Add**:
  - **Active Directory Domain:** Choose the Active Directory domain associated with the keytab file.
  - **DNS Principal:** The name and domain of the DNS server receiving the DDNS updates. Note that this value is not the same as the Kerberos principal you specified when you generated the keytab file.  
Use the following format when you complete this field: **DNS/dns\_server\_fqdn@ad\_domain**  
*dns\_server\_fqdn:* This is the FQDN of the DNS server. You can use the "dig" command to perform a DNS lookup to obtain the FQDN of the DNS server as it appears on the SOA record.  
*ad\_domain:* This is the Active Directory domain of the DNS server.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

### Verifying the Configuration

After you configure the AD domain controller and the Infoblox DHCP server, you can view the syslog of the Infoblox DHCP server to verify if it successfully established a security context with the AD domain controller. The DHCP server displays a series of messages similar to the following:

```
dhcpcd: Enabled GSS-TSIG for zone corpxyz. using principal jdoe/anywhere@corpxyz.LOCAL
.
dhcpcd: GSS-TSIG security thread has started.
dhcpcd: GSS-TSIG security update starting at 1222389338.
dhcpcd: Acquiring GSS-TSIG credential for jdoe/anywhere@corpxyz.LOCAL.
dhcpcd: Acquired GSS-TSIG credential for jdoe/anywhere@corpxyz.LOCAL(good for 3568s).
dhcpcd: Security context established with server 10.34.123.4 for principal[ jdoe/
anywhere@corpxyz.LOCAL|mailto:jdoe/anywhere@corpxyz.LOCAL] (good for 568s).
dhcpcd: GSS-TSIG security update complete at 1222389338. Next update in 360s.
```

In addition, you can log in to the Infoblox CLI and use the `show dhcp_gss_tsig` CLI command to troubleshoot your configuration. For information about this command, refer to the *Infoblox CLI Guide*.

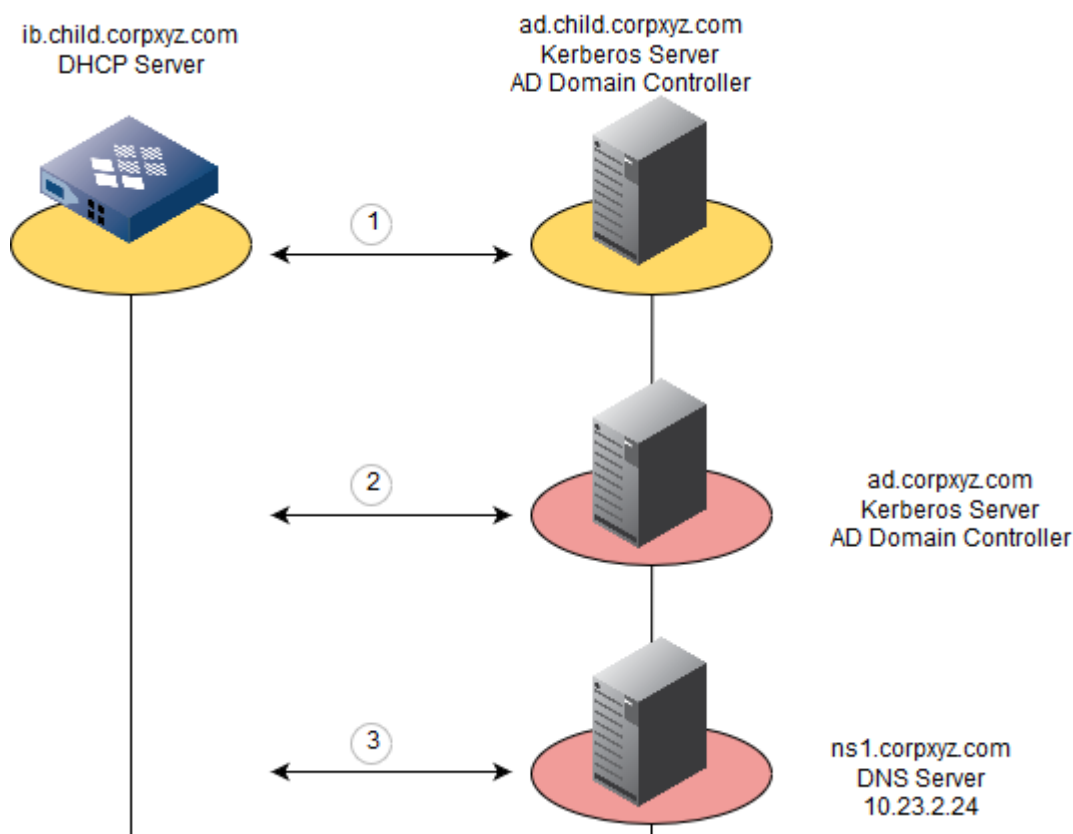
### Sending Secure DDNS Updates to a DNS Server in Another Domain

Domain and forest trust relationships provide clients authenticated access to resources in other domains. Some trusts are automatically created, such as the two-way, direct trust between parent and child domains in a forest. Other trusts must be created manually. Refer to the Microsoft Active Directory documentation for information on establishing trusts between domains.

Once a direct trust exists between two AD domains, a KDC from one domain can grant a referral to the KDC of the other domain. The Infoblox DHCP server can then use the referral to request access to services in the other domain.

In the following figure the Infoblox DHCP server in the child.corpxyz.com domain needs to send GSS-TSIG authenticated DDNS updates to the DNS server in its parent domain, corpxyz.com domain. There is an automatic two-way trust between the domains because corpxyz.com domain is the parent of child.corpxyz.com domain.

#### *Sending Secure DDNS Updates to a DNS Server in Another Domain*



After you configure the Infoblox DHCP server and AD domain controller, the following occurs:

1. Kerberos – In Same Domain  
The Infoblox DHCP server uses the TGT (ticket-granting ticket) from the AD/Kerberos server, ad.child.corpxyz.com, to request a service ticket for DNS/ns1.corpxyz.com@corpxyz.COM. The Kerberos server replies with a referral ticket for the Kerberos server in the corpxyz.com domain, ad.corpxyz.com.
2. Kerberos — In the Other Domain  
The Infoblox DHCP server uses the referral ticket and requests a service ticket from ad.corpxyz.com for DNS/ns1.corpxyz.com@corpxyz.COM. The Kerberos server replies with a service ticket for DNS/ns1.corpxyz.com@corpxyz.COM.
3. TKEY Negotiations (GSS Handshake)  
The Infoblox DHCP server sends the DNS server ns1.corpxyz.com a TKEY (transaction key) request, which includes the service ticket. The DNS server replies with a TKEY response that includes a TSIG (transaction signature). The Infoblox appliance and the DNS server have established a security context, enabling the DHCP server to send DDNS updates to the DNS server.

### Configuring DHCP to Send GSS-TSIG Updates to Another Domain

Before the DHCP server can send secure DDNS updates to a DNS server in a different domain, you must ensure that a direct trust relationship exists between the domain of the DHCP server and that of the DNS server. (For information, refer to the Active Directory documentation.)

Following are the tasks to configure the AD domain controller and the Infoblox DHCP server for secure updates to another domain. All the configuration is done on the AD domain controller for the domain of the DHCP server and on the Infoblox DHCP server.:

1. Complete the following tasks on the AD domain controller for the domain of the DHCP server:
  - a. Add a user account for the Infoblox DHCP server. In the configuration example, the user account is ibdhcp. For information, see [Creating an AD User Account](#) above.

- b. Generate the keytab file for the Infoblox DHCP server and export it from the AD domain controller to a local directory on your management system. For the DHCP server in the Sending Secure DDNS Updates to a DNS Server in Another Domain, the principal is `ibdhcp/ib.child.corpxyz.com@CHILD.corpxyz.COM`. For information, see [Generating and Exporting the Keytab File](#) above.
2. Complete the following tasks on the Infoblox DHCP server:
  - a. Import the keytab file from your management system to the appliance and enable GSS-TSIG dynamic updates at the Grid or member level. For information, see [Enabling GSS-TSIG Authentication for DHCP](#).
  - b. Configure the external forward-mapping zone for the DDNS updates. Note that the DNS principal uses the domain of the DNS server, regardless of the domain of the DHCP server. For the DNS server in the figure Sending Secure DDNS Updates to a DNS Server in Another Domain above, the DNS principal is `DNS/ns1.corpxyz.com@corpxyz.COM`. For information, see [Managing GSS-TSIG keys](#).

### Configuration Example

Following are the steps to configure the example shown in the figure Sending Secure DDNS Updates to a DNS Server in Another Domain above:

On the Active Directory domain controller:

1. Create a user account for the Infoblox DHCP server. The user account is **ibdhcp**.

2. Generate the keytab file and export it to your management system. If the domain controller is running Windows Server 2003:

```
ktpass -princ ibdhcp/ib.child.corpxyz.com@CHILD.corpxyz.COM -mapuser ibdhcp@CHILD.corpxyz.COM -pass infoblox -out ibdhcp.ktb -ptype krb5_nt_principal -crypto des-cbc-md5 +desonly
```

On the Infoblox DHCP server:

1. Enable GSS-TSIG at the member level.
2. From the **DHCP** tab, click the **Members** tab -> *member* checkbox -> Edit icon.
3. On the **DDNS** -> **Basic** tab of the editor, complete the following:
  - **Override**: Select this checkbox.
  - **DDNS Updates**: Select the **Enable DDNS Updates** checkbox.
  - **GSS-TSIG**: Select **Override**, and then complete the following:
    - **Enable GSS-TSIG Updates**: Select this checkbox.
    - **Domain Controller (KDC)**: Enter `ad.child.corpxyz.com`. This is the KDC in the domain of the DHCP server.
    - **Manage Key tab Files**: Click **Manage Key tab Files**. In the *Manage GSS-TSIG Keys* dialog box, click the Add icon. Click **Select**, navigate to the keytab file, select the keytab file that you just uploaded, `ibdhcp/ib.child.corpxyz.com@CHILD.corpxyz.COM`, and then click **Upload**. Click **Close**.
4. Save the configuration and click **Restart** if it appears at the top of the screen.
5. Configure the external forward mapping zone, `corpxyz.com`:
  - a. On the **DHCP** tab, expand the Toolbar, and then click **Configure DDNS**.
  - b. In the *DNS Updates to External Zones* table of the *DDNS Properties* editor, click the Add icon and complete the following fields in the *Add External DDNS Zone* panel:
    - **Zone Name**: Enter `corpxyz.com`.
    - **DNS Server Address**: Enter the IP address of the primary DNS server to which the Infoblox DHCP server sends DDNS updates. In the example, the DNS server is `ns.corpxyz.com`. Therefore, enter its IP address, which is `10.23.2.24`.
    - **Security**: Choose **GSS-TSIG** from the drop-down list, complete the following, and then click **Add**:
      - **Active Directory Domain**: Choose `child.corpxyz.com`.
      - **DNS Principal**: Enter `DNS/ns1.corpxyz.com@corpxyz.COM`.
6. Save the configuration and click **Restart** if it appears at the top of the screen.

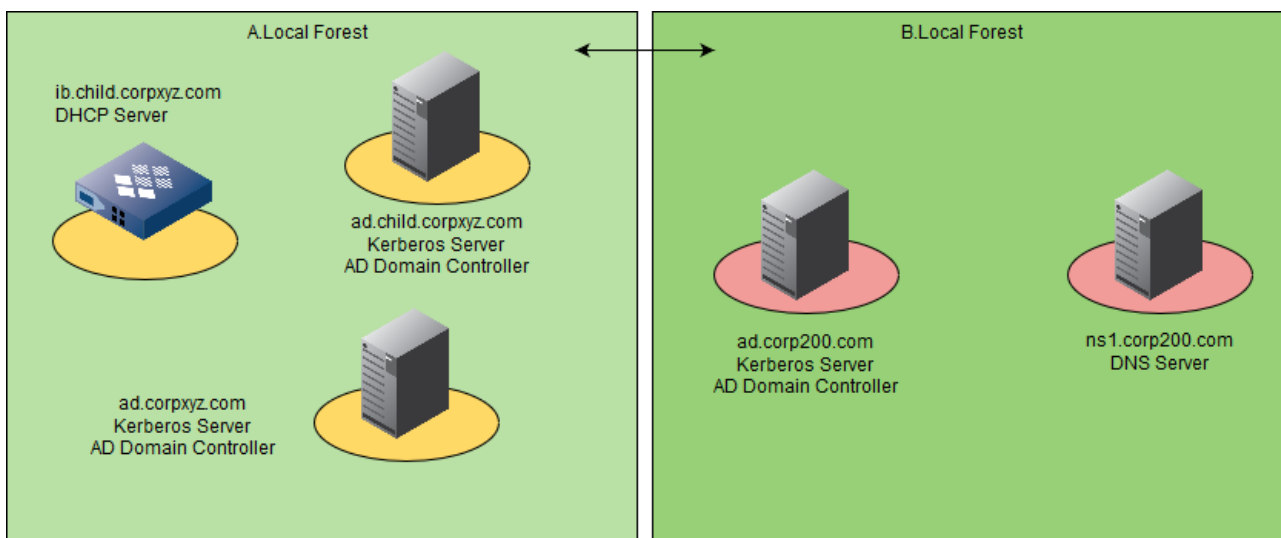
## Sending GSS-TSIG Updates to a DNS Server in Another Forest

The Infoblox DHCP server can also send secure DDNS updates to a DNS server that belongs to a domain in another forest, as long as a forest trust exists. Refer to the Microsoft Active Directory documentation for information on establishing forest trusts.

Similar to the authentication process between domains, the authentication process between forests also uses referrals. The appliance follows the referral chain until it reaches the domain controller of the domain in which the service is located. Note that forest trusts are not transitive. For example, if the DHCP server is in forest A and the DNS server is in forest C, a direct trust must exist between forest A and forest C for the DDNS updates to succeed. Having a trust between forest A and B, and between forest B and C is not sufficient.

In the below figure, a trust exists between the A.Local forest and the B.Local forest. The Infoblox DHCP server in the A.Local forest needs to dynamically update the DNS server in the B.Local forest.

### *Sending Secure DDNS Updates to a DNS Server in Another Forest*



The following authentication process occurs:

1. Kerberos – In Same Domain  
The Infoblox appliance uses the TGT (ticket-granting ticket) from the AD/Kerberos server, ad.child.corpxyz.com, to request a service ticket for DNS/ns1.corp200.com@CORP200.COM. The Kerberos server does not find the principal name in its domain database and after consulting the global catalog, it replies with a referral ticket for its parent domain.
2. Kerberos — Referral Chain  
The appliance contacts a domain controller in corpxyz.com and requests a referral to a domain controller in the corp200.com domain in B.Local Forest. When it receives the referral, the DHCP server contacts the domain controller and requests a service ticket for the DNS server, ns1.corp200.com. The domain controller replies with a service ticket for [ DNS/ ns1.corp200.com@CORP200.COM. |mailto:DNS/ns1.corp200.com@CORP200.COM]
3. TKEY Negotiations (GSS Handshake)  
The Infoblox appliance sends the DNS server ns1.corp200.com a TKEY (transaction key) request, which includes the service ticket. The DNS server replies with a TKEY response that includes a TSIG (transaction signature). The Infoblox appliance and the DNS server have established a security context.

## Configuring DHCP to Send GSS-TSIG Updates to a Different Forest

Configuring the Infoblox DHCP server for dynamic updates to a DNS server in another forest is similar to the configuration used to send dynamic updates to another domain in the same forest. For information, see [Configuring DHCP to Send GSS-TSIG Updates to Another Domain](#) section.

## Configuring GSS-TSIG keys

You can upload keytab files that contain a single GSS-TSIG key or multiple GSS-TSIG keys on a single NIOS appliance. For each member in the Grid, you can upload up to 256 GSS-TSIG keys in a single keytab file. Trust relationships between AD domains and AD forests are not required. You can upload GSS-TSIG keys through Grid Manager or the Infoblox API.

Note that only superusers can manage all GSS-TSIG keys globally on a given member through Grid Manager or the Infoblox API. Using this feature, superusers can determine the keys that belong to a particular member. You can assign multiple GSS-TSIG keys to a member and all these keys are saved in the Grid. The uploaded keys will be available in the member DNS, Grid DNS, member DHCP or Grid DHCP properties. NIOS supports the following GSS-TSIG encryption types:

- des-cbc-crc
- des-cbc-md5
- arcfour-hmac-md5
- aes128-cts-hmac-sha1-96
- aes256-cts-hmac-sha1-96

NIOS displays a warning message in Grid Manager and in the syslog if you upload a key that does not belong to the GSS-TSIG encryption types. For more information, see Logging Messages below.

### Limitations when Using Multiple GSS-TSIG keys

- You can assign SPNs belonging to different domains to a DNS member, but you cannot assign SPNs belonging to different domains to a DHCP member, although two DHCP members can update the same DNS member.
- You must ensure that the domains assigned to a DNS member are unique.
- The GSS-TSIG domain for a remote forward or remote reverse zone is single-valued. For example, if DHCP clients ABC and XYZ from Grid 1 want to send DDNS updates to Grid 2, either client ABC or XYZ will succeed.

### Scheduled Upgrade

A scheduled upgrade with one or more keys in the keytab files that you have uploaded will operate the same as prior to upgrade. NIOS will parse and extract keys from the uploaded keytab file. NIOS automatically assigns these keys to the DNS member, DHCP member, Grid DHCP or Grid DNS to which the keytab file was uploaded before the upgrade. You can assign these keys to Grid members after the upgrade is complete.

NIOS does not display an error message if the keys do not have an SPN with the DNS prefix, but it will record a warning message in the syslog.

### Admin Permissions for Configuring GSS-TSIG keys

You can assign a key to a Grid member only if you have read permission for the kerberos key and read/write permission for the member. You can upload keys only if you have read/write permissions for kerberos keys. To remove a key that is assigned to a member, you must have read/write permission for the respective member.

Note that in the **Administration** -> **Administrators** -> **Permissions** tab, NIOS displays **All Kerberos Keys** and **Kerberos Key** in the **Resource** and **Resource Type** columns respectively for **DHCP Admin** and **DNS Admin** roles with default read/write permissions.

### Enabling GSS-TSIG Authentication for DHCP

You can enable GSS-TSIG authentication at the Grid or member level and associate it with one or more keys of the same SPN or realm. When you enable GSS-TSIG authentication, make sure that you upload the keytab file from the Kerberos account for the Infoblox DHCP server. You can import keytab files with multiple keys to the Grid or to individual members. You can assign the uploaded keys to member DHCP or Grid DHCP. The appliance displays a warning message if you assign a GSS-TSIG key with service class "DNS" in its SPN to a DHCP member. The appliance displays an error message in the following cases:

- if you assign keys of different realms to a DHCP member or Grid DHCP.

- when you try to enable GSS-TSIG without a valid key.

The AD domain controller stores the keytab file in the directory in which you generated the keytab file. You can copy this file to a management system that connects to the NIOS appliance or launch the NIOS Grid Manager on the AD domain controller and import the keytab file to the NIOS appliance.

To enable GSS-TSIG authentication for DHCP and import keytab files:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Properties**.  
**Member:** From the **Data Management** tab, select the **DHCP** tab and click the **Members** tab -> *member* checkbox -> Edit icon. To override an inherited property, click **Override** next to it and complete the appropriate fields.  
**Standalone DHCP:** From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **System DHCP Properties**.
2. In the **IPv4 DDNS** -> **Basic** tab or the **IPv6 DDNS** -> **Basic** tab of the editor, complete the following:
  - **DDNS Updates:** Select **Enable DDNS Updates** to enable the DHCP servers in the Grid to send DDNS updates.
  - **DDNS Domain Name:** Specify the domain name of the network that the appliance uses to update DNS. For IPv4 clients, you can specify this at the network, network template, range, and range template levels. For IPv6 clients, you can specify this at the Grid, member, network, shared network, and network template levels.
  - **DDNS Update TTL:** You can set the TTL used for A record and PTR records updated by the DHCP server. The default is shown as zero. If you do not enter a value here, the appliance by default sets the TTL to half of the DHCP lease time with a maximum of 3600 seconds. For example, a lease time of 1800 seconds results in a TTL of 900 seconds, and a lease time of 86400 seconds results in a TTL of 3600 seconds.
  - **DDNS Update Method:** Select the method used by the DHCP server to send DDNS updates. You can select either **Interim** or **Standard** from the drop-down list. The default is **Interim**. When you select **Interim**, TXT record will be created for DDNS updates and when you select **Standard**, DHCID record will be created for DDNS updates. But in the **IPv4 DDNS** -> **Advanced** tab or the **IPv6 DDNS** -> **Advanced** tab, if you have selected **No TXT Record** mode for the DHCP server to use when handling DNS updates, then TXT record or DHCID record is not created for DDNS updates.  
 If you change the DDNS update method from **Interim** to **Standard** or vice versa, then the DHCP server changes the DHCID type used from TXT record to DHCID record or vice versa as the leases are renewed.  
 This is supported for clients that acquire both IPv4 and IPv6 leases. Infoblox recommends you to configure different DDNS update method for IPV4 leases and IPV6 leases, **Interim** for IPv4 lease and **Standard** for IPv6 lease.
  - **GSS-TSIG:** Complete the following:
    - **Enable GSS-TSIG Updates:** Select this to enable the DHCP server to send GSS-TSIG authenticated DDNS updates. The DHCP server uses the KDC server configured in the member properties. Therefore, the **Enable GSS-TSIG Updates** field must be overridden in the *Member DHCP Properties* dialog box to use the KDC configured in the *Member DHCP Properties* dialog box.
    - **Manage Keytab Files:** To upload a keytab file, click **Manage GSS-TSIG keys**. In the *Manage GSS-TSIG Keys* dialog box, click the Add icon. In the *Upload* dialog box, click **Select**, navigate to the keytab file, select it, and then click **Upload**. You can also delete individual keys. For more information about managing GSS-TSIG keys, see *Managing GSS-TSIG keys* below.
    - **Domain Controller:** Enter the resolvable host name or IP address of the AD domain controller that hosts the KDC for the domain.
    - **Principal:** The principal member of the key. For GSS-TSIG based DDNS updates, the SPN of the key used to carry out the update does not require the server class 'DHCP.' You can either specify an FQDN or an IP address for the <host> of an SPN.
    - **GSS-TSIG Key:** Select the name of the GSS-TSIG key from the drop-down list that you want the Grid to use. This is only available if you have uploaded a keytab file. Click the arrow beside the Add icon to either assign keys or upload and assign keys. You can either select **Assign Keys** or **Upload & Assign Keys** from the drop-down list.
      - **Assign Keys:** Select **Assign Keys** to select a GSS-TSIG key from the *GSS-TSIG Key Selector*. Click **Principal**, which is displayed as a hyperlink, to select it. For more information about the **GSS-TSIG Key Selector**, see *Selecting Keys in the GSS-TSIG Key Selector* below.



- **Upload & Assign Keys:** Select **Upload & Assign Keys** to upload and assign keys. In the *Upload* dialog box, select the file and navigate to the file you want to upload. Click **Upload**. The appliance assigns the keys contained in the selected keytab file.
  - The following are displayed in the table:
    - **Version:** The version of the key.
    - **Encryption type:** The encryption type of the key.
    - **Last update:** The timestamp when the key was uploaded.
  - **Zones this member can update securely:** Click **Display** to list the external zones to which the Grid member can send secured DDNS updates.
  - **Lease Renewal Update:** Select **Update DNS on DHCP Lease Renewal** to enable the DHCP server to update DNS when a DHCP lease is renewed.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Deleting GSS-TSIG keys associated with DHCP Objects

You can delete individual keys if it is not in use by the Grid or any member. To delete a key that is assigned to a member, you must have Read/Write permission for the member. To delete individual keys:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **Grid DHCP Properties**.  
**Member:** From the **Data Management** tab, select the **DHCP** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
**Standalone DHCP:** From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **System DHCP Properties**.
2. In the **IPv4 DDNS** tab or the **IPv6 DDNS** -> **Basic** tab of the editor, select keys from the list under **GSS-TSIG Keys** and click the Delete icon to delete keys.

## Enabling GSS-TSIG Authentication for DNS

For GSS-TSIG based DDNS updates, the SPN of the key used to carry out the update must have 'DNS' in its service class. You can upload a keytab file to the Grid with multiple keys in which each key has an SPN in this format: DNS/<host>@<realm>. You can associate a DNS member or a Grid DNS with one or more keys of the same SPN or realm or of different SPN or realms. You can assign the uploaded keys to member DNS or Grid DNS, but NIOS displays an error when you try to enable GSS-TSIG without a valid key if the assigned key does not have the service class 'DNS' in its SPN.

To enable GSS-TSIG authentication for DNS and import keytab files:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon. To override an inherited property, click **Override** next to it and complete the appropriate fields.  
**Standalone DNS:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **System DNS Properties**.
2. In the **GSS-TSIG** -> **Basic** tab of the editor, complete the following:
  - **GSS-TSIG:** Select **Enable GSS-TSIG authentication of clients** to accept GSS-TSIG signed DDNS updates from clients that belong to different AD domains in which each domain has a unique GSS-TSIG key.
  - **Manage Keytab Files:** To upload a keytab file, click **Manage GSS-TSIG keys**. In the *Manage GSS-TSIG Keys* dialog box, click the Add icon. In the *Upload* dialog box, click **Select**, navigate to the keytab file, select it, and then click **Upload**. You can also delete individual keys. For more information, see *Managing GSS-TSIG keys* below.
  - **GSS-TSIG Keys:** Click the arrow beside the Add icon to either assign keys or upload and assign keys. You can either select **Assign Keys** or **Upload & Assign Keys** from the drop-down list.
    - **Assign Keys:** Select **Assign Keys** to select a GSS-TSIG key from the *GSS-TSIG Key Selector*. Click **Principal**, which is displayed as a hyperlink, to select it. For more information about the *GSS-TSIG Key Selector*, see *Selecting Keys in the GSS-TSIG Key Selector* below.
    - **Upload & Assign Keys:** Select **Upload & Assign Keys** to upload and assign keys. In the *Upload* dialog box, select the file and navigate to the file you want to upload. Click **Upload**. The appliance assigns keys in the uploaded file. The following are displayed:
    - **Principal:** The principal member of the key. For GSS-TSIG based DDNS updates, the SPN of the key used to carry out the update must have DNS in its service class. It is of the following form:

```
DNS/<host>@<realm>
```

You can either specify an FQDN or an IP address for the <host> of an SPN.

- **Domain:** The domain name assigned to the DNS member.
- **Version:** The version of the key.
- **Encryption type:** The encryption type of the key.
- **Last update:** The timestamp when the key was uploaded.

3. Save the configuration.

NIOS sorts the data in the table based on the last updated timestamp, by default. Note that sometimes GSS-TSIG updates might stop working after you restart the DNS service because the appliance discards the GSS-TSIG keys, when you restart the DNS service. If this happens, wait several minutes until the Microsoft server performs another handshake using the new key.

### Deleting GSS-TSIG keys associated with the DNS Objects

You can delete individual keys if it is not in use by the Grid or any member. To delete a key that is assigned to a member, you must have Read/Write permission for the member. To delete individual keys:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.  
**Standalone DNS:** From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **System DNS Properties**.
2. In the **GSS-TSIG** -> **Basic** tab of the editor, select keys from the list under **GSS-TSIG Keys**.
3. Click the Delete icon to delete.

### Logging Messages

The appliance saves the audit log entries for insert and delete operations. If you upload keys with encryption types other than the ones that NIOS supports, the appliance displays a warning message in Grid Manager and in the syslog and also it displays the encryption type as \*other\* in Grid Manager and in the syslog. For more information about the syslog, see [Using a Syslog Server](#).

The appliance generates an audit log when you upload a key, assign the key to a member, remove the key associated with a member or delete a key. The audit log entries are based on each key that you have uploaded. For example, NIOS saves the following in the audit log when you upload a key:

```
2014-02-14 18:17:30.531Z \[admin\]: imported DNS Kerberos key for
principal='DNS/infoblox.localdomain@abc.com', version=5, enctype=des-cbc-crc
```

For more information about audit logs, see [Using the Audit Log](#). You can search Kerberos keys using the realm (domain), principal name or an encryption type.

The appliance generates a comment in the option section of the DNS configuration file for each Kerberos principal that is associated with the Grid member. These comments are for information only and it indicates the principals, their versions and encryption types that are used by the appliance.

### Managing GSS-TSIG keys

You can upload a keytab file that contains one or multiple GSS-TSIG keys and delete multiple keys through the *Manage GSS-TSIG Keys* wizard. To manage multiple GSS-TSIG keys, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, expand the Toolbar and click **Manage GSS-TSIG Keys**.
2. In the *Manage GSS-TSIG Keys* wizard, the following are displayed:



- **Principal:** The principal name that is mapped to the keytab file.
  - **Domain:** The name of the domain that is mapped to the keytab file.
  - **Version:** The version of the keytab file.
  - **In use:** Indicates whether the keytab file is in use or not.
  - **Members:** The members associated with the keytab file. Click the hyperlink and the *Members* dialog box is displayed. It displays the list of members that are associated with the keytab file.
  - **Encryption type:** The encryption type of the key.
  - **Last update:** The timestamp when the key was last uploaded.
3. Click the Upload Keytab File icon to upload a new keytab file. In the *Upload* dialog box, click **Select** and navigate to the keytab file. Click **Upload** to upload the file.

To delete a GSS-TSIG key, select the appropriate key and click the Delete icon.

### Selecting Keys in the GSS-TSIG Key Selector

NIOS displays the keys that you have uploaded using the keytab files. You can choose a filter and an operator to view specific keys that you have uploaded. The *GSS-TSIG Key Selector* wizard is displayed only when you select **Assign Keys** in the *Properties* editor. For more information about how to assign keys to DNS and DHCP objects, see *Enabling GSS-TSIG Authentication for DNS and Enabling GSS-TSIG Authentication for DHCP* above respectively.

To select a key from the *GSS-TSIG Key Selector*, complete the following:

1. Click **Show Filter** to filter the values:
  - Select a value from the drop-down list to filter your values: **Domain**, **Encryption type**, **In use**, **Last update**, **Principal**, and **Version**.
  - Select one of these operators from the drop-down list: **equals**, **does not equal**, **begins with**, and **does not begin with**.
  - Enter the value that you want to search in the text box. Click **Hide Filter** to hide the filter. Alternatively, you can enter a value in the text box for **Find** and click **Go** to search specific keys from the keytab files.
2. The following details are displayed in the table:
  - **Principal:** The principal name that is mapped to the keytab file. Click **Principal** to assign the key to the DNS or DHCP object.
  - **Domain:** The name of the domain that is mapped to the keytab file.
  - **Version:** The version of the keytab file.
  - **In use:** Indicates whether the keytab file is in use or not.
  - **Members:** The members associated with the keytab file.
  - **Encryption type:** The encryption type of the key.
  - **Last update:** The timestamp when the key was last uploaded.

### Accepting DDNS Updates from DHCP Clients

A NIOS appliance serving DNS can support Active Directory and accept both unauthenticated and GSS-TSIG authenticated updates from DHCP clients, DHCP servers, and AD domain controllers. The appliance supports servers running Windows 2000 Server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, or Windows Server 2016 with the Active Directory service installed. When adding a NIOS appliance that serves DNS to an AD environment, you must configure the AD/Kerberos server and NIOS appliance as follows—based on whether or not you want the DNS server to support DDNS updates using GSS-TSIG authentication:

- AD/Kerberos Server
  1. Enable zone transfers to the NIOS appliance.
  2. (For GSS-TSIG) Create a user account for the NIOS appliance that it can use for authentication.
  3. (For GSS-TSIG) Generate the keytab file of the DNS server and save it to your management system.
- NIOS Appliance
  4. (GSS-TSIG) Enable GSS-TSIG support.

5. (GSS-TSIG) Import the keytab file of the DNS server from your management system to the NIOS appliance.
6. (GSS-TSIG) Enable GSS-TSIG authentication.
7. Add a forward-mapping zone and give it a name matching the AD DNS zone whose resource records you want to import.
8. Specify the domain controller from which the appliance can receive DDNS updates. An AD domain controller replicates its data among other domain controllers within its AD domain and among domain controllers in other domains.
9. Import zone data from the specified domain controller.
10. Enable the acceptance of DDNS updates from the AD domain controller and from the DHCP clients and servers whose addresses the DHCP server assigns. You can set this at the Grid, member, and zone levels.
11. (For GSS-TSIG) Enable acceptance of GSS-TSIG DDNS updates from the AD domain controller and from the addresses that the DHCP server assigns. You can set this at the Grid, member, and zone levels.

As you can see from the above task list, adding a NIOS appliance that serves DNS to an AD environment without GSS-TSIG support involves four simple steps. To include GSS-TSIG support, there are several additional steps.

### Supporting Active Directory and Unauthenticated DDNS Updates

Before configuring the NIOS appliance, configure the AD domain controller to permit zone transfers to the IP address of the appliance. Then on the appliance, you can do the following to configure a forward-mapping zone to support AD (Active Directory) and receive unauthenticated DDNS updates from DHCP clients, DHCP servers, and AD domain controllers.

- Create a forward-mapping zone, as described in [Creating an Authoritative Forward-Mapping Zone](#). Give it a name that matches the AD DNS zone whose resource records you want to import.
- Specify the domain controllers from which the appliance can receive updates, as described in [Configuring AD Support](#) below.
- Import the zone data from the domain controller. For information, see [Importing Data into Zones](#).
- Enable the appliance to accept DDNS updates from the DHCP clients and servers whose addresses the DHCP server assigns. You can set this at the Grid, member, and zone levels. For information, see [Enabling DNS Servers to Accept DDNS Updates](#).

### Configuring AD Support

You can configure a forward-mapping zone to support AD from the *Active Directory* wizard or from the **Active Directory** tab of the *Authoritative Zone* editor. This section describes both methods.

To configure AD support using the *Active Directory* wizard:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Configure Active Directory**. Note that from the **Zones** tab, you must select a zone before you click **Configure Active Directory**.
2. In the *Active Directory* wizard, complete the following, and then click **Next**:
  - **Select Zone**: Click this and select a zone. The name of the zone must match the name in the AD domain controller so the zone transfer from the AD domain controller to the NIOS appliance can succeed.
  - **Allow unsigned updates from Domain Controllers**: Select this option. If you have configured DNS resolvers in the Grid, the appliance sends DNS queries for the names and addresses of the AD domain's domain controllers. Since the name of the zone that you selected is the same as the AD domain name on the domain controller, the appliance can then send a DNS query for the SRV records attached to the domain name. It also sends a DNS query for the A record of each domain controller to determine its IP address. The query results are listed in the next panel.
3. You can edit the list of domain controllers, if necessary. Click **Next** to proceed to the next step.
  - To add a domain controller, click the Add icon and specify the IP address.
  - To delete a domain controller from the list, select it and click the Delete icon.
4. Complete the following:
  - **Do you want to create underscore zones to hold the records added by the Domain Controllers?**

This option allows the appliance to create the following subzones that the DNS server must have to answer AD-related DNS queries:

`_msdcs.zone`  
`_sites.zone`  
`_tcp.zone`  
`_udp.zone`  
`domaindnszones.zone`  
`forestdnszones.zone`

Note that these zones are automatically generated. You cannot edit these zones or import data into them. They cannot be modified, thus providing protection against forged updates.

5. Save the configuration and click **Restart** if it appears at the top of the screen.

To configure AD support using the *Authoritative Zone* editor:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox -> Edit icon.
2. In the *Authoritative Zone* editor, select the **Active Directory** tab and do the following:
  - **Allow unsigned updates from these Domain Controllers:** Select this checkbox and specify the AD domain controllers from which the appliance can receive DDNS updates.
  - **Automatically create underscore zones:** Select this checkbox to automatically create the subzones.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

You can then import zone data, as described in [Importing Data into Zones](#).

## Accepting GSS-TSIG-Authenticated Updates

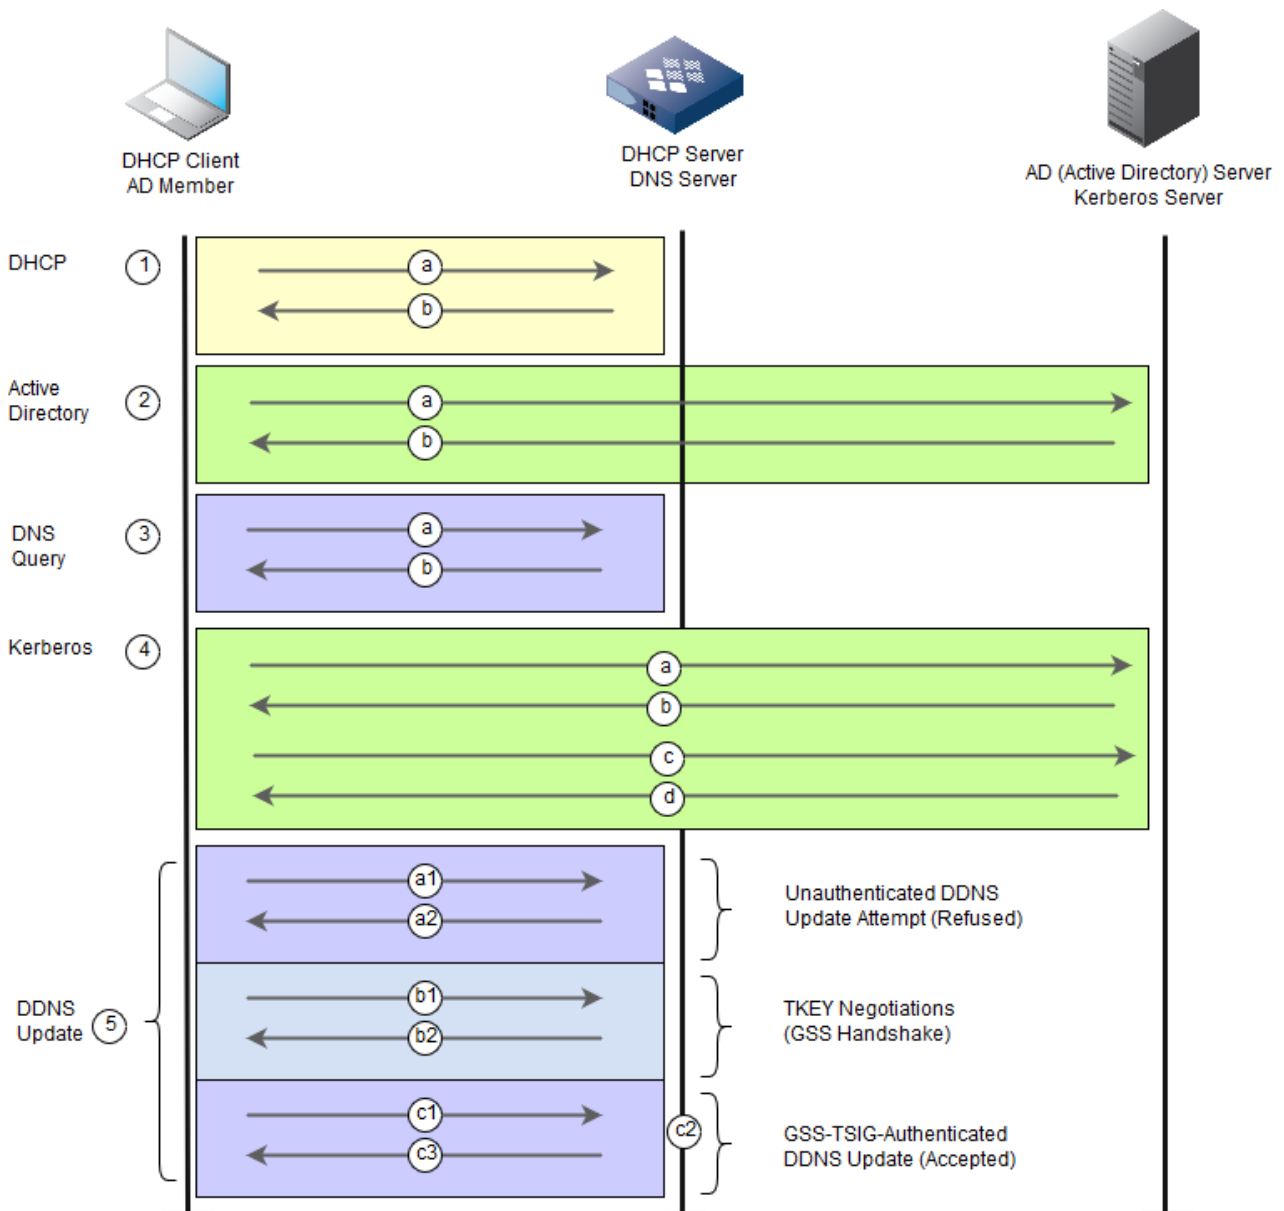
A NIOS appliance can support Active Directory and process secure GSS-TSIG-authenticated DDNS updates from DHCP clients, DHCP servers, and AD domain controllers. The appliance supports servers running Windows 2000 Server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, or Windows Server 2016 with the Active Directory service installed. The process in which a DHCP client dynamically updates its resource records on a DNS server using GSS-TSIG authentication is shown in *Figure 21.10* below. The illustration also shows the relationship of the clients, the DHCP server, the DNS server, and the Kerberos server (running on the AD domain controller).



### Note

For explanations of the alphanumerically notated steps in *Figure 21.10* below, see the section following the illustration.

*Figure 21.10 Authenticating DDNS Updates with GSS-TSIG*



1. DHCP – IP Address and Network Parameters Assignment
  - a. The DHCP client requests an IP address.
  - b. The DHCP server assigns an IP address, subnet mask, gateway address, and a DNS server address.
2. Active Directory – Computer and User Logins
  - a. The computer sends a DNS request to locate the AD domain controller, and then logs in to the domain controller.  
 Note that the computer accounts have passwords that the AD domain controller and computer maintain automatically. There are two passwords for each computer: a computer account password and a private key password. By default, both passwords are automatically changed every 30 days.
  - b. The user manually logs in to a domain.
3. DNS – Query for the Kerberos Server
  - a. The computer (or *client*) automatically sends a query for *kerberos.\_udp.dc.\_msdcs.dom\_name\_* to the DNS server whose IP address it received through DHCP.
  - b. The NIOS appliance replies with the name of the Kerberos server.
4. Kerberos – Login, and TGT and Service Ticket Assignments
  - a. The client automatically logs in to the Kerberos server.
  - b. The Kerberos server sends the client a TGT (ticket-granting ticket).

- c. Using the TGT, the AD member requests a service ticket for the DNS server.
  - d. The Kerberos server replies with a service ticket for that server.
5. DDNS – Dynamic Update of the Client's Resource Records
- a. Unauthenticated DDNS Update Attempt (Refused)
    - i. The client sends an unauthenticated DDNS update.
    - ii. The DNS server refuses the update.
  - b. TKEY negotiations (GSS Handshake):
    - i. The client sends the DNS server a TKEY (transaction key) request. A Transaction Key record establishes shared secret keys for use with the TSIG resource record. For more information, see *RFC 2930, Secret Key Establishment for DNS (TKEYRR)*. The request includes the service ticket. The service ticket includes the appliance's principal and proposed TSIG (transaction signature) key, along with other items such as a ticket lifetime and a timestamp.
    - ii. The DNS server responds with a DNS server-signed TSIG, which is a "meta-record" that is never cached and never appears in zone data. A TSIG record is a signature of the update using an HMAC-MD5 hash that provides transaction-level authentication. For more information, see *RFC 2845, Secret Key Transaction Authentication for DNS (TSIG)*. The two participants have established a security context.
  - c. GSS-TSIG-Authenticated DDNS Update (Accepted).
    - i. The client sends an authenticated DDNS update, which includes the following resource records:
      - A – Address record
      - or
      - PTR – Pointer record
      - TKEY – Transaction Key record
      - TSIG – TSIG record
    - ii. The DNS server authenticates the DDNS update and processes it.
    - iii. The DNS server sends a GSS-TSIG-authenticated response to the AD member, confirming the update.



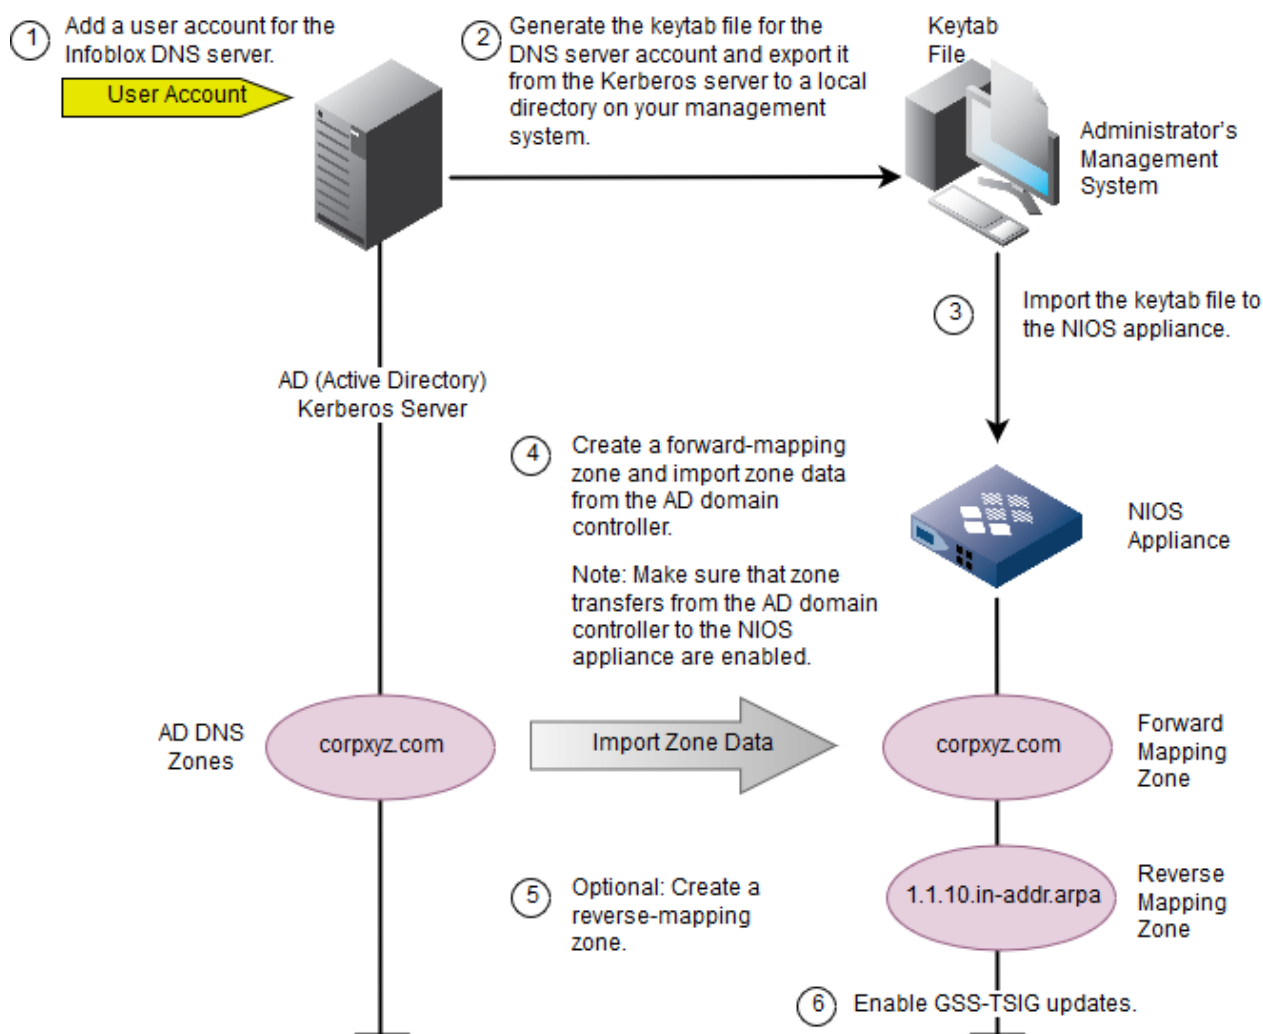
#### Note

For GSS-TSIG authentication to work properly, the system clock times of the Infoblox DHCP server, AD domain controller and DNS server must be synchronized. One approach is to use NTP and synchronize all three devices with the same NTP servers.

## Configuring DNS to Receive GSS-TSIG Updates

You can configure an appliance to support Active Directory and accept secure DDNS updates from clients using GSS-TSIG. The initial configuration tasks are shown in *Figure 21.11*. The appliance supports servers running Windows 2000 Server, Windows Server 2003, Windows Server 2008, Windows Server 2008 R2, Windows Server 2012, Windows Server 2012 R2, or Windows Server 2016 with the Active Directory service installed.

*Figure 21.11 Adding a NIOS Appliance to an AD Environment with GSS-TSIG Support*



On an already functioning AD domain controller:

1. Enable zone transfers to the NIOS appliance.
2. Add a user account for the NIOS appliance serving DNS. A corresponding account on the Kerberos server is automatically created. For information, see [Creating an AD User Account](#) below.
3. Export the keytab file for the NIOS appliance account from the Kerberos server to a local directory on your management system. For information, see [Generating and Exporting the Keytab File](#) below.

On an Infoblox appliance:

1. Import the keytab file from your management system to the Infoblox appliance and enable GSS-TSIG authentication on the appliance. For information, see [Importing the Keytab File](#) and [Enabling GSS-TSIG Authentication](#) below.
2. Configure a forward-mapping zone with the same name as the AD zone. For information, see [Creating an Authoritative Forward-Mapping Zone](#).
3. (Optional) Create a reverse-mapping zone for the network address space that corresponds to the domain name space in the forward-mapping zone. For information, see [Creating an Authoritative Reverse-Mapping Zone](#).
4. Import the zone data from the AD domain controller. For information, see [Importing Zone Data](#).
5. Enable the acceptance of GSS-TSIG-signed updates from the AD controller and from the DHCP clients and servers whose addresses the DHCP server assigns. For information, see [Accepting GSS-TSIG Updates](#) below.

## Creating an AD User Account

Connect to the AD domain controller and create a user account for the NIOS appliance.



### Note

The name you enter in the User logon name is the name that you later use when exporting the keytab file. This is also the principal name. The text in the First name, Initials, Last name, and Full name fields is irrelevant to this task.

The AD domain controller automatically creates a Kerberos account for this user with an accompanying keytab. Note the following:

- If you define an expiration date for the user account and you later create a new account when the first one expires, the keytab for the corresponding Kerberos account changes. At that point, you must update the keytab file on the NIOS appliance (see [Generating and Exporting the Keytab File](#) and [Importing the Keytab File and Enabling GSS-TSIG Authentication](#) below). Optionally, if your security policy allows it, you can set the user account for the NIOS appliance so that it never expires.
- If the AD domain controller is running Windows Server 2003, the user account must have the DES encryption type enabled. You can enable this either in the Account tab when you create the user account or by specifying **+DesOnly** when you use the Ktpass tool to generate the keytab file.
- The newly created AD user account must be a member of the DnsUpdateProxy group or an account that allows it to update records that have potentially been added by another DHCP server, such as DNS Admins.

## Generating and Exporting the Keytab File

You can generate and export the keytab file for the Kerberos account by using the Ktpass tool. Note that the version of the Ktpass tool that you use must match the Windows version of the domain controller. For example, if you are using a domain controller running Windows Server 2008 or Windows Server 2008 R2, you must use the Ktpass tool for Windows Server 2008 or Windows Server 2008 R2.

You enter different commands for generating and exporting the keytab file, depending on whether you are generating the keytab file from a server running Microsoft Windows 2000, Windows Server 2003, Windows Server 2008, or Windows Server 2008 R2.

### Generating the Keytab on Windows 2000

To export the keytab file using a Microsoft Windows 2000 Resource Kit:

1. Start a command prompt.
2. Enter the following command to export the keytab file for the NIOS appliance user account:

```
C:> ktpass -princ service_name/FQDN_instance@REALM -mapuser AD_username -pass  
password -out  
filename.keytab
```

For example:

```
C:> ktpass -princ DNS/ns1.corpxyz.com@corpxyz.COM -mapuser ns1@corpxyz.com -pass  
37Le37  
-out ns1.keytab
```

### Generating the Keytab on Windows Server 2003

The Ktpass tool is included in the Windows Server 2003 Support Tools. To export the keytab file using a Microsoft Windows 2003 Resource Kit:

1. Start a command prompt.
2. Enter the following command to export the keytab file for the NIOS appliance user account:

```
ktpass -princ DNS/FQDN_instance@REALM -mapuser AD_username -pass password -out filename.keytab -ptype KRB5_NT_PRINCIPAL -crypto des-cbc-md5 +DesOnly
```

For example:

```
ktpass -princ DNS/ns1.corpxyz.com@corpxyz.COM -mapuser ns1@corpxyz.com -pass 37Le37 -out ns1.keytab -ptype KRB5_NT_PRINCIPAL -crypto des-cbc-md5 +DesOnly
```

where:

**-princ** = Kerberos principal

- DNS = Service name in uppercase format
- ns1.corpxyz.com = Instance in FQDN (fully-qualified domain name) format; this is the same as the DNS name of the NIOS appliance
- corpxyz.COM = The Kerberos realm in uppercase format; this must be the same as the AD domain name

**-mapuser** = Maps the Kerberos principal name to the AD user account

- ns1@corpxyz.com = The AD user name for the NIOS appliance

**-pass** = The AD user account password

- 37Le37 = The password of the user account for the NIOS appliance

**-out** = Exports the keytab file

- ns1.keytab = The name of the keytab file

**-ptype** = Sets the principal type. This must be **krb5\_nt\_principal**.

**-crypto** = Specifies the encryption type. This must be **des-cbc-md5**.

**+DesOnly** = Specifies DES encryption for the account. Include this if you did not enable DES encryption for the account.

### Generating the Keytab on Windows Server 2008/Windows Server 2008 R2

A Windows Server 2008 or Windows Server 2008 R2 domain controller allows you to generate a keytab file with multiple keys for one principal. The Infoblox DNS server accepts GSS-TSIG updates from DHCP clients that provide a Kerberos ticket for any of the keys in its configured keytab. To generate the keytab file using the Ktpass tool:

1. Start a command prompt.
2. Enter the following command to export the keytab file for the NIOS appliance user account:

```
ktpass -princ DNS/FQDN_instance@REALM -mapuser AD_username -pass password -out filename.keytab -ptype krb5_nt_principal -crypto encryption
```

For example:



```
ktpass-princDNS/ns1.corpxyz.com@corpxyz.COM-mapusersns1@corpxyz.com-pass37Le37-  
outns1.keytab-ptype krb5_nt_principal-cryptoAES256-SHA1
```

where:

**-princ** = Kerberos principal

- DNS = Service name in uppercase format
- ns1.corpxyz.com = Instance in FQDN format; this is the same as the DNS name of the NIOS appliance
- corpxyz.COM = The Kerberos realm in uppercase; this must be the same as the AD domain name

**-mapuser** = Maps the Kerberos principal name to the AD user account

- ns1@corpxyz.com = The AD user name for the NIOS appliance

**-pass** = The AD user account password

- 37Le37 = The password of the user account for the NIOS appliance

**-out** = Exports the keytab file

- ns1.keytab = The name of the keytab file

**-ptype** = Sets the principal type. This must be **krb5\_nt\_principal**.

**-crypto** = Specifies the encryption type.

After you execute the command to generate the keytab file, the AD domain controller displays a series of messages similar to the following to confirm that it successfully generated the keytab file:

```
Targeting domain controller: qacert.test.local Using legacy password setting method  
  
Successfully mapped DNS/ns1.corpxyz.com to ns1.  
  
Key created.  
  
Output keytab to ns1.keytab:  
  
Keytab version: 0x502  
  
keysize 80 DNS/ns1.corpxyz.com@corpxyz.COM ptype 1 (KRB5_NT_PRINCIPAL) vno 3 etype  
0x12 (AES256-SHA1)  
  
keylength 32 (0xea8675d7abf13fd760a744088642fb917ceb6c9d267f5c54e595597846f06407)
```



#### Note

The keytab file contains highly sensitive data for the NIOS appliance account. Ensure that you store and transport its contents securely.

## Modifying an AD User Account

To change any AD user account information (login, password, etc):

1. Remove the previous user account from AD.
2. Create a new user for GSS-TSIG mapping.
3. Generate a new keytab file.
4. Import the keytab file to the DNS server.

## Importing the Keytab File and Enabling GSS-TSIG Authentication

Before you can enable GSS-TSIG authentication, you must import the keytab file from the Kerberos account for the NIOS appliance. To import the keytab file:

1. From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.
2. In the *Member DNS Properties* editor, click **Toggle Expert Mode**.
3. When the additional tabs appear, click **GSS-TSIG** and do the following:  
If a principal name and version number are listed, there is a keytab file loaded on the appliance. Compare this information with that for the NIOS appliance account on the Kerberos server to make sure that they match. If there is no keytab file on the NIOS appliance or if the loaded keytab file does not match that on the Kerberos server, you must load the correct keytab file.
  - Click **Upload**, click **Browse** to navigate to the keytab file, and then click **Upload**.
  - **Enable GSS-TSIG authentication of clients**: Select this checkbox.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

Each time you export a keytab file from a Kerberos server running on Windows Server 2003, the version number of the keytab file increases incrementally. Because the version number on the keytab file that you import to the NIOS appliance must match the version that is in use on the Kerberos server, you should select the last keytab file that is exported from the Kerberos server if you have exported multiple keytab files. (A Kerberos server running on Windows 2000 does not increase the version number of keytab files with each export.)

## Accepting GSS-TSIG Updates

You can allow a Grid or specific members or zones to accept GSS-TSIG signed updates from domain controllers and DHCP clients and servers, as follows:

1. **Grid**: From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.  
**Member**: From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox -> Edit icon.  
**Zone**: From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox -> Edit icon.  
To override an inherited property, click **Override** next to it and complete the appropriate fields.
2. Select the **Updates** tab and do the following in the **Basic** subtab:
  - **Allow GSS-TSIG signed updates**: Select this option.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

You can then use the *Active Directory* wizard or navigate to the **Active Directory** tab of the *Authoritative Zone* editor to enable the appliance to create underscore zones for the records hosted by domain controllers and to allow GSS-TSIG signed updates to the underscore zones.

To use the *Active Directory* wizard:

1. From the **Data Management** tab, select the **DNS** tab, expand the Toolbar and click **Configure Active Directory**.
2. In the *Configure Active Directory* wizard, complete the following, and then click **Next**:
  - **Select Zone**: Click this and select a zone. The name of the zone must match the name in the AD domain controller so the zone transfer from the AD domain controller to the NIOS appliance can succeed.
  - **Allow GSS-TSIG-signed (secure) updates from Domain Controllers**: Select this option.
3. Complete the following:
  - **Do you want to create underscore zones to hold the records added by the Domain Controllers?**  
This option allows the appliance to create the following subzones that the DNS server must have to answer AD-related DNS queries:
    - \_msdcs.zone*
    - \_sites.zone*
    - \_tcp.zone*
    - \_udp.zone*
    - domaindnszones.zone*
    - forestdnszones.zone*Note that these zones are automatically generated. You cannot edit these zones or import data into them.
  - **Allow GSS-TSIG-signed updates to underscore zones**: Select this checkbox to allow underscore zones to accept GSS-TSIG signed updates.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

To use the *Authoritative Zone* editor:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox -> Edit icon.
2. In the *Authoritative Zone* editor, select the **Active Directory** tab and do the following:
  - **Allow unsigned updates from these Domain Controllers:** Clear this checkbox.
  - **Automatically create underscore zones:** (select)  
This option automatically creates the following subzones that the DNS server must have to answer AD-related DNS queries:
    - \_msdcs.zone*
    - \_sites.zone*
    - \_tcp.zone*
    - \_udp.zone*
    - domaindnszones.zone*
    - forestdnszones.zone*Note that these zones are automatically generated and cannot be manually edited.
  - **Allow GSS-TSIG-signed updates to underscore zones:** Select this checkbox to allow underscore zones to accept GSS-TSIG signed updates.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Secure Dynamic Updates

The secure dynamic updates feature provides several methods to restrict dynamic DNS updates to certain records. This includes:

- Restrictions for updates to statics records. For more information, see [Restricting Updates to Static Records below](#).
- Restrictions for updates to records marked as protected. For more information, see [Restricting Updates to Protected Records below](#).
- Restrictions based on GSS-TSIG principal authentication. For more information, see [Restricting Updates Based on GSS-TSIG Principal Authentication below](#).
- Restrictions based on FQDN patterns. For more information, see [Restricting Updates Based on FQDN Patterns below](#).

Only static and dynamic record source type support secure dynamic updates. You can see the record source type in the Resource Record Viewer. The following table shows which type of secure dynamic updates is applicable to different record source types.

### *Secure Dynamic Update Types*

Secure Dynamic Update Type	Record Source Type
Restrictions for updates to statics records	Static
Restrictions for updates to protected records	Static, dynamic
Restrictions based on GSS-TSIG principal authentication	Dynamic
Restrictions based on FQDN patterns	Static, dynamic

Sometimes when the updating record has the same data as the existing record, you may need to initialize the record creation timestamp to avoid unwanted DNS record scavenging. For more information, see [Forcing Creation Timestamp Initialization for Unchanged Records below](#).

Failed attempts to dynamically update secured records are recorded in the NIOS syslog. You can view it, as described in [Viewing the Syslog](#) and [Searching in the Syslog](#).

You can use Smart Folders to organize data by record source, principal, or protection state. For more information, see [Smart Folders](#).

In addition, you can use Global Search to search for records by principal name. For more information, see [Using Global Search](#).



#### Note

To use the secure dynamic updates feature, you must have a DNS license installed in the Grid Manager.

## Restricting Updates to Static Records

This method prevents updates to all RRsets containing static records at once in the Grid, DNS view, or zone. To prevent updates to specific static records, see [Restricting Updates to Protected Records](#) below.



#### Note

When you upgrade from a previous NIOS version to NIOS 7.3 or later, all dynamic updated records are labelled as static records if you enable the Secure Dynamic Updates feature. Infoblox suggests that you enable this feature only after all records are changed to Dynamic. NIOS tags the RRsets that are not auto-generated as static records.

To restrict updates to all static records in the Grid, DNS view, or zone:

1. In the Grid DNS, view, or zone properties, click **Updates -> Advanced**.
2. To override the inherited properties, click **Override**.
3. Under **Secure Dynamic Updates**, select **Prevent dynamic updates to RRsets containing static records**.
4. Click **Save & Close**.

## Restricting Updates to Protected Records

You can restrict updates to the records of your choice, by marking them as protected. You can do this for both static and dynamic records. The Resource Record Viewer displays the protection status of the records in the Protected column: Yes or No.

You can protect the following record types:

- A record
- AAAA record
- CNAME record
- DNAME record
- MX record
- NAPTR record
- PTR record
- SRV record
- TXT record
- Host record

For all the above mentioned records except the host record, you can change the type from static to dynamic and back, if required.

To restrict updates to protected records:

1. In the DNS Resource Records viewer, select a record or multiple records.
2. In the Toolbar, select **Protect Records -> Enable Protection**.  
Or  
In the properties dialog for a record, click **Updates**, select the **Protected** checkbox, and then click **Save & Close**.
3. Enable updates prevention at the corresponding level:
  - a. In the Grid DNS, view, or zone properties, click **Updates -> Advanced**.
  - b. If necessary, click **Override** to override the inherited properties.

- c. Select **Prevent dynamic updates to RRsets containing protected records**.
- d. Click **Save & Close**.

## Restricting Updates Based on GSS-TSIG Principal Authentication

This method implies tracking the Kerberos GSS-TSIG principal that created a record and restricting DDNS updates attempted by a different GSS-TSIG principal on this record.

The Resource Record Viewer displays the GSS-TSIG authentication information in the Principal column: it displays the principal name if the client that created the record is authenticated and the principal is tracked.

The tracked principal is also displayed in the record properties. You can change the principal associated to a record by clicking **Select Principal** in the record properties and specifying the required principal.

Additionally, you can use dynamic update groups to manage the allowed principals. For more information, see [About Dynamic Update Groups](#) below.

To restrict updates based on GSS-TSIG principal authentication:

1. In the Grid DNS, view, or zone properties, click **Updates** -> **Advanced**.
2. To override the inherited properties, click **Override**.
3. Under **Secure Dynamic Updates**, select **Track the GSS-TSIG principals that create dynamic records**.  
Note that for this option to work, ensure that you have selected **Enable GSS-TSIG authentication of clients** in the GSS-TSIG properties of the Grid or the corresponding zone or view.
4. Select **Require the appropriate GSS-TSIG principal to update RRsets that track principals**.
5. Optionally, specify an active dynamic update group.
6. Click **Save & Close**.

## About Dynamic Update Groups

In some cases, for example, in DHCP failover associations, you need to allow different GSS-TSIG principals to update each other's records. To that end, you can join multiple principals into clusters, where all principals are considered as equivalent and therefore can update affected records without being their originators. You can join multiple clusters into a dynamic update group. The clusters within a group, however, are not considered equivalent and cannot update each other's records.

When you have several dynamic update groups defined, you can assign different groups to be active for the Grid, a DNS view, or a zone as described in [Restricting Updates Based on GSS-TSIG Principal Authentication](#) below. If no group is assigned, then no principals are considered to be equivalent.

For information on how to add dynamic update groups and clusters, see [Managing Dynamic Update Groups and Clusters](#) below.



### Note

Viewing and modifying the configuration of a dynamic update group requires Grid DNS permissions. Selecting a group as active for the Grid, a view, or a zone requires read permission on the Grid DNS, as well as write permission on the object being modified.

## Managing Dynamic Update Groups and Clusters

To add a dynamic update group:

1. In **Data Management** -> **DNS**, expand the Toolbar and click **Manage Dynamic Update Groups**.
2. Click the Add icon.
3. Select **Add Dynamic Update Group**.
4. Specify the group name.
5. Optionally, provide a comment.
6. Click **Save and Close**. Proceed to adding clusters to the group as described below.

To add a cluster:

1. In the *Manage Dynamic Update Groups* window, click the Add icon.
2. Select **Add Cluster**.
3. Select the dynamic update group in which you want to include the cluster.

4. Specify the cluster name.
5. Optionally, provide a comment.
6. Click **Save and Close**.
7. To add principals to the cluster, select the cluster in the *Manage Dynamic Update Groups* window and click the Add icon. A principal can appear in multiple clusters.
8. Select one of the following:
  - **Add Principal**: This adds a new row in the table. Specify the principal name in the row.
  - **Select Principal**: This opens the *Principal Selector* dialog. Select the required principal from the list.
9. Click **Close**.

To edit or delete a group, cluster, or principal, select it in the *Manage Dynamic Update Groups* window, and click the corresponding icon.

You can also export data about dynamic update groups, their clusters, and principals in the Infoblox CSV Import format by clicking the Export icon in the *Manage Dynamic Update Groups* window. For more information, see [Exporting Data to Files](#).

### Restricting Updates Based on FQDN Patterns

This implies defining FQDN patterns for domain names which prevents DDNS updates to matching FQDNs. To restrict updates based on FQDN patterns:

1. In the Grid DNS, view, or zone properties, click **Updates -> Advanced**.
2. To override the inherited properties, click **Override**.
3. Under **Dynamic Update Patterns**, select **Prevent dynamic updates to FQDNs matching these patterns** and specify patterns:
  - To add an FQDN pattern, click the Add icon and specify a pattern in the new table row. Not that to use the DNS Traffic Control LBDN wild cards to specify FQDN patterns. For more information, see [Configuring LBDN Patterns](#).
  - To delete an FQDN pattern, select the checkbox next to the pattern and click the Delete icon.
4. Click **Save & Close**.

### Forcing Creation Timestamp Initialization for Unchanged Records

If the attributes of a resource record do not change in the result of a DDNS update, Grid Manager ignores the update and the record's creation timestamp remains the same. This may cause valid records with outdated timestamp to be removed during DNS scavenging. To avoid this, you can set the record creation time to be modified even when the record data do not change at DDNS update. You can do this for the whole Grid, or for a specific DNS view or authoritative zone.

To force the creation timestamp initialization for unchanged resource records:

1. Open the *Grid DNS Properties*, *DNS View Properties*, or *Authoritative Zone Properties* editor.
2. For a DNS view or authoritative zone, click **Override**.
3. Select **Modify creation time even when resource record data is unchanged**.
4. Click **Restart** in the Grid Manager's system messages banner for the setting to take effect.

For information about DNS scavenging, see [DNS Record Scavenging](#).

## Configuring DNSSEC

This section provides general information about DNSSEC. The topics in this section include:

- [DNSSEC](#)
- [DNSSEC Resource Records](#)
- [DNSKEY Resource Records](#)
- [RRSIG Resource Records](#)
- [NSEC/NSEC3 Resource Records](#)
- [NSEC3PARAM Resource Records](#)
- [DS Resource Records](#)

- [Configuring DNSSEC on a Grid](#)
- [Enabling DNSSEC](#)
- [Setting DNSSEC Parameters](#)
- [Signing a Zone](#)
- [About HSM Signing](#)
- [Configuring Grid Members to Support DNSSEC as Secondary Servers](#)
- [Configuring Recursion and Validation for Signed Zones](#)
- [Applying Policies and Rules to DNS Queries that Request DNSSEC Data](#)

---

## DNSSEC

DNSSEC (DNS Security Extensions) provides mechanisms for authenticating the source of DNS data and ensuring its integrity. It protects DNS data from certain attacks, such as man-in-the-middle attacks and cache poisoning. A man-in-the-middle attack occurs when an attacker intercepts responses to queries and inserts false records. Cache poisoning can occur when a client accepts maliciously created data. DNSSEC helps you avoid such attacks on your networks. DNSSEC provides changes to the DNS protocol and additional resource records (RRs) as described in the following RFCs:

- *RFC 4033, DNS Security Introduction and Requirements*
- *RFC 4034, Resource Records for the DNS Security Extensions*
- *RFC 4035, DNSSEC Protocol Modifications*
- *RFC 4641, DNSSEC Operational Practices*
- *RFC 4956, DNS Security (DNSSEC) Opt-In*
- *RFC 4986, Requirements Related to DNS Security (DNSSEC) Trust Anchor Rollover*
- *RFC 5155, DNS Security (DNSSEC) Hashed Authenticated Denial of Existence*
- *RFC 5702, Use of SHA-2 Algorithms with RSA in DNSKEY and RRSIG Resource Records for DNSSEC*

DNSSEC uses public key cryptography to authenticate the source of DNS responses and to ensure that DNS responses were not modified during transit. Public key cryptography uses an asymmetric key algorithm. With asymmetric keys, one key is used to decrypt data that was encrypted using the other key.

In DNSSEC, the primary name server of a zone generates at least one public/private key pair. It "signs" each data set in the zone by running it through a one-way hash, and then encrypting the hash value with the private key. The public key is stored in an RR type introduced by DNSSEC, the DNSKEY RR. Resolvers use the DNSKEY record to decrypt the hash value. If the hash values match, then the resolver is assured of the authenticity of the message.

In addition to the DNSKEY record, DNSSEC also introduces new RRs which DNS servers can use to authenticate the non-existence of servers, zones, or resource records. For information about the DNSSEC resource records, see [DNSSEC Resource Records](#).

DNSSEC uses the EDNS0 message extension. Resolvers include the EDNS OPT pseudo-RR with the DO (DNSSEC OK) bit set to indicate that they are requesting DNSSEC records. A DNS client or resolver sets the EDNS DO bit when it sends a query for data in a signed zone. When the DNS server receives such a query, it includes the additional DNSSEC records in its response, according to the DNSSEC standard rules. In addition, because DNSSEC messages are often large, the EDNS0 message extension also provides mechanisms for handling larger DNS UDP messages. For information about EDNS0, refer to *RFC 2671, Extension Mechanisms for DNS (EDNS0)*. For information about the DO bit, refer to *RFC 3225, Indicating Resolver Support of DNSSEC*.



### Warning

*When you disable EDNS0 on the appliance, all outgoing DNSSEC queries to zones within trusted anchors will fail even if DNSSEC validation is enabled. To ensure that DNSSEC functions properly, do not disable EDNS0 on the appliance. For more information, see [Using Extension Mechanisms for DNS \(EDNS0\)](#).*

DNSSEC also supports new data in the packet header, the CD (Checking Disabled) bit and the AD (Authenticated Data) bit. The CD bit is used by resolvers in their DNS queries and the AD bit is used by recursive name servers in their responses to queries.

A resolver can set the CD bit in its query to indicate that the name server should not validate the DNS response and that the resolver takes responsibility for validating the DNS data it receives.

A name server that has successfully validated the data in a DNS response sets the AD (Authenticated Data) bit in the message header to indicate that all resource records in its response have been validated and are authentic. Note that unless the connection between the DNS server and client has been secured, such as through TSIG, the client cannot rely on the AD bit to indicate valid data. The data could have been changed in transit between the server and client.

Resolvers can trust a response with the AD bit set only if their communication channel is secure.

You can also configure the NIOS appliance to always apply RPZ policies, DNS blacklists, or NXDOMAIN rules to DNS responses, regardless of whether the queries request DNSSEC data. For more information about how to configure this, see [Applying Policies and Rules to DNS Queries that Request DNSSEC Data](#). For information about RPZ policies, DNS blacklists, and NXDOMAIN rules, see their respective sections in this guide.

Related topic

[Configuring DNSSEC](#)

## DNSSEC Resource Records

Following are the DNSSEC RR types:

- DNS Public Key (DNSKEY) resource records — For information, see [DNSKEY Resource Records](#).
- Resource Record Signature (RRSIG) records — For information, see [RRSIG Resource Records](#).
- Next Secure (NSEC/NSEC3) records — For information, see [NSEC/NSEC3 Resource Records](#).
- NSEC3PARAM records — For information, see [NSEC3PARAM Resource Records](#).
- Delegation Signer (DS) resource records — For information, see [DS Resource Records](#).

For detailed information about each RR, refer to *RFC 4034, Resource Records for the DNS Security Extensions* and *RFC 5155, DNS Security (DNSSEC) Hashed Authenticated Denial of Existence*.



### Note

The appliance supports IDNs for DNSKEY records, DS records, NSEC records, NSEC3PARAM records, and RRSIG records.

## DNSKEY Resource Records

When an authoritative name server digitally signs a zone, it typically generates two key pairs, a zone-signing key (ZSK) pair and a key-signing key (KSK) pair. The name server uses the private key of the ZSK pair to sign each RRset in a zone (an RRset is a group of resource records that are of the same owner, class, and type.). It stores the public key of the ZSK pair in a DNSKEY record. The name server then uses the private key of the KSK pair to sign all DNSKEY records, including its own, and stores the corresponding public key in another DNSKEY record. As a result, a zone typically has two DNSKEY records; a DNSKEY record that holds the public key of the ZSK pair, and another DNSKEY record for the public key of the KSK pair.



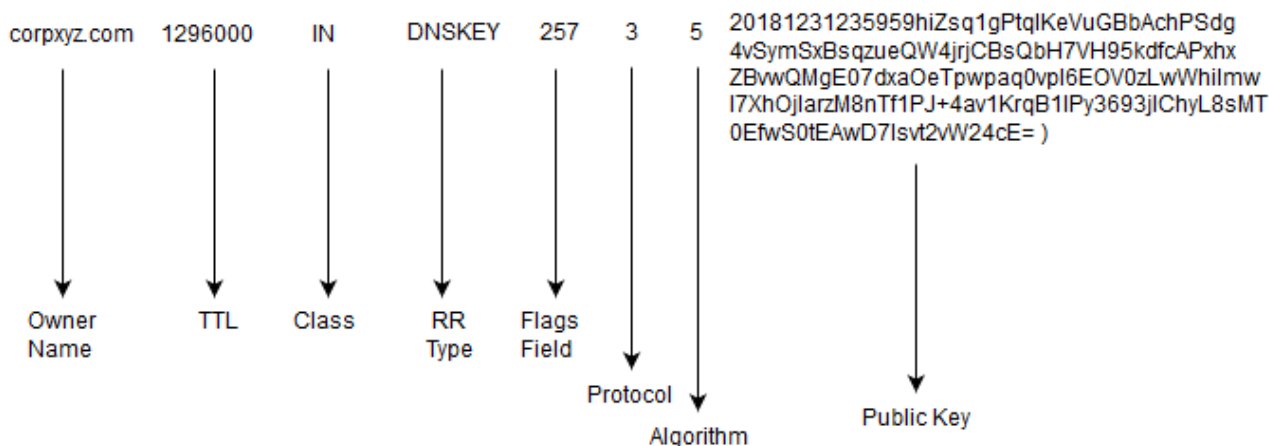
### Note

For the remainder of this chapter, the DNSKEY record that holds the public key of the ZSK pair is referred to as the ZSK and the DNSKEY record that holds the public key of the KSK pair is referred to as the KSK.



The purpose of the KSK is two-fold. First, it is referenced in the Delegation Signer (DS) RR that is stored in a parent zone. The DS record is used to authenticate the KSK of the child zone, so a resolver can establish a chain of trust from the parent zone to its child zone. For more information about the DS RR, see [DS Resource Records](#). Second, if a zone does not have a chain of trust from a parent zone, security aware resolvers can configure the KSK as a trust anchor; that is, the starting point from which it can build a chain of trust from that zone to its child zones. Note that though the two key pairs, KSK and ZSK, are used in most DNSSEC environments, their use is not required by the RFCs. A zone administrator can use a single private/public key pair to sign all zone data. (Note that Infoblox appliances require two key pairs.)

Following is an example of a DNSKEY RR:



The first four fields specify the domain name of the zone that owns the key, the resource record TTL, class, and RR type. The succeeding fields are:

- **Flags Field:** In its wire format, this field is two bytes long. (The wire format is used in DNS queries and responses.) Bits 0 through 6 and 8 through 14 are reserved, and have a value of 0. Bit 7 indicates if the record holds a DNS zone key. Bit 15 is the Secure Entry Point (SEP) flag, which serves as a hint that indicates whether the DNSKEY record contains a ZSK or a KSK, as described in *RFC 3757, DNSKEY RR SEP Flag*. Zone administrators typically set the SEP flag of a DNSKEY record of a zone when it contains the KSK, to indicate that it can be used as a trust anchor. However, a DNSKEY record that does not have the SEP flag set can also be used as a trust anchor.

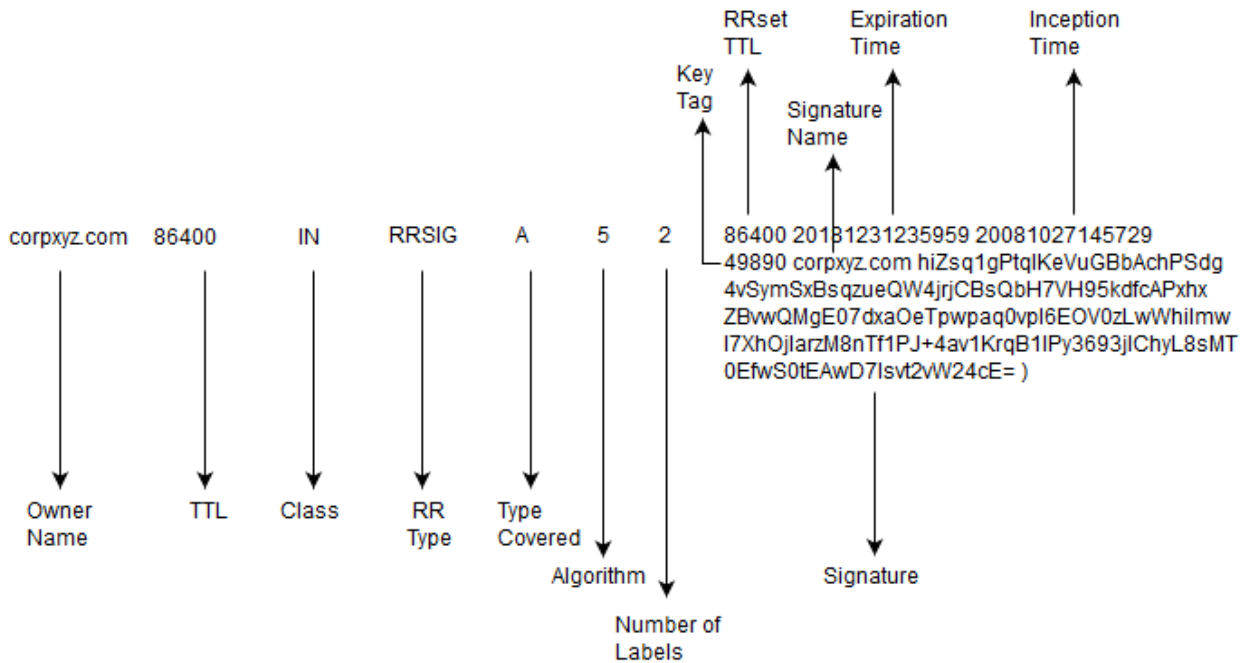
Given the currently defined flags, in its text format, the flags field is represented as an unsigned decimal integer with the possible values of 0, 256 and 257. A value of 256 indicates that the DNSKEY record holds the ZSK and a value of 257 indicates that it contains the KSK. In general, this field contains an odd number when the DNSKEY record holds the KSK.

- **Protocol:** This always has a value of 3, for DNSSEC.
- **Algorithm:** Identifies the cryptographic algorithm of the public key. The available types are:
  - 1 = RSA/MD5
  - 2 = Diffie-Hellman (This is not supported by BIND and Infoblox appliances.)
  - 3 = DSA
  - 4 = Reserved
  - 5 = RSA/SHA1
  - 6 = DSA/SHA1/NSEC3
  - 7 = RSA/SHA1/NSEC3
  - 8 = RSA/SHA-256
  - 10 = RSA/SHA-512
  - 13 = ECDSAP/SHA-256
  - 14 = ECDSAP/SHA-384
- **Public Key:** The public key encoded in Base64.

## RRSIG Resource Records

A signed zone has multiple RRsets, one for each record type and owner name. (The owner is the Fully Qualified Domain Name of the RRset.) When an authoritative name server uses the private key of the ZSK pair to sign each RRset in a zone, the digital signature on each RRset is stored in an RRSIG record. Therefore, a signed zone contains an RRSIG record for each RRset.

Following is an example of an RRSIG record:




The first four fields specify the owner name, TTL, class, and RR type. The succeeding fields are:

- Type Covered: The RR type covered by the RRSIG record. The RRSIG record in the example covers the A records for corpxyz.com.
- Algorithm: The cryptographic algorithm that was used to create the signature. It uses the same algorithm types as the DNSKEY record indicated in the Key Tag field.
- Number of Labels: Indicates the number of labels in the owner name of the signed records. There are two labels in the example, corpxyz and com.
- RRset TTL: The TTL value of the RRset covered by the RRSIG record.
- Expiration Time: The signature expiration time in UTC format.
- Inception Time: The signature inception time in UTC format.
- Key Tag: The key tag value of the DNSKEY RR that validates the signature.
- Signature Name: The zone name of the RRset.
- Public Key: The Base64 encoding of the signature.

## NSEC/NSEC3 Resource Records

When a name server receives a request for a domain name that does not exist in a zone, the name server sends an authenticated negative response in the form of an NSEC or NSEC3 RR. NSEC and NSEC3 records contain the next secure domain name in a zone and list the RR types present at the NSEC or NSEC3 RR's owner name. The difference between an NSEC and NSEC3 RRs is that the owner name in an NSEC3 RR is a cryptographic hash of the original owner name prepended to the name of the zone. NSEC3 RRs protect against zone enumeration.

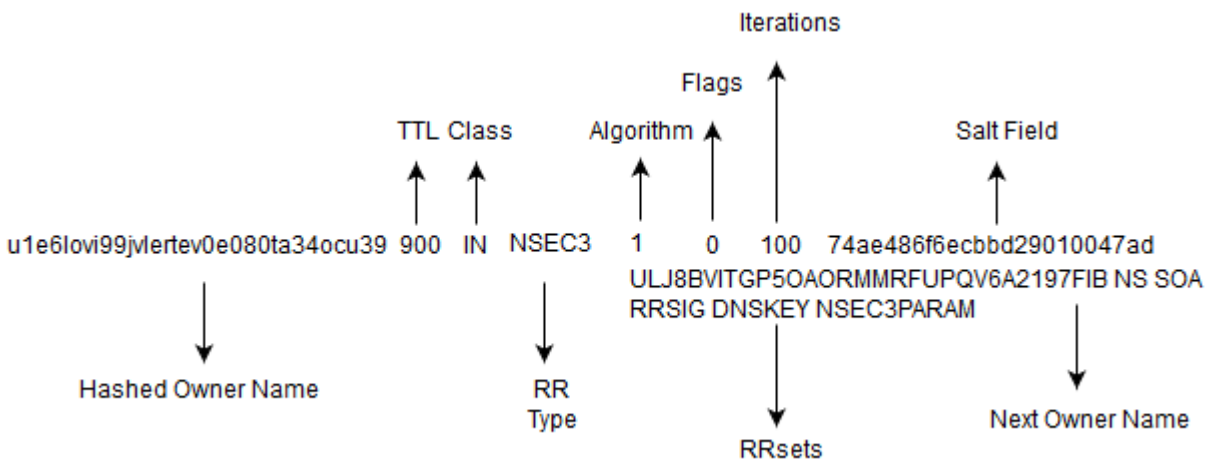
Following is an example of an NSEC record:



The first four fields specify the owner name, TT, class and RR type. The succeeding fields are:

- Next Owner Name: In the canonical order of the zone, the next owner name that has authoritative data or that contains a delegation point NS record.
- RRsets: The RRsets that exist at the owner name of the NSEC record, which are NS, SOA, RRSIG, NSEC, and DNSKEY in the example.

Following is an example of an NSEC3 RR:




The first field contains the hashed owner name. It is followed by the TTL ,class and RR type. The fields after the RR type are:

- Algorithm: The hash algorithm that was used. The currently supported algorithm is SHA-1, which is represented by a value of 1.
- Flags Field: Contains 8 one-bit flags, of which only one flag, the Opt-Out flag, is defined by RFC 5155. The Opt-Out flag indicates whether the NSEC3 record covers unsigned delegations.
- Iterations: The number of times the hash function was performed.
- Salt Field: A series of case-insensitive hexadecimal digits. It is appended to the original owner name as protection against pre-calculated dictionary attacks.
- Next Owner Name: Displays the next hashed owner name.
- RRsets: The RR types that are at the owner name.

## NSEC3PARAM Resource Records

An authoritative DNS server uses NSEC3PARAM RRs to determine which NSEC3 records it includes in its negative responses. An NSEC3PARAM RR contains the parameters that an authoritative server needs to calculate hashed owner names. As stated in RFC 5155, the presence of an NSEC3PARAM RR at a zone apex indicates that the specified parameters may be used by authoritative servers to choose an appropriate set of NSEC3 RRs for negative responses.

Following is an example of an NSEC3PARAM record:



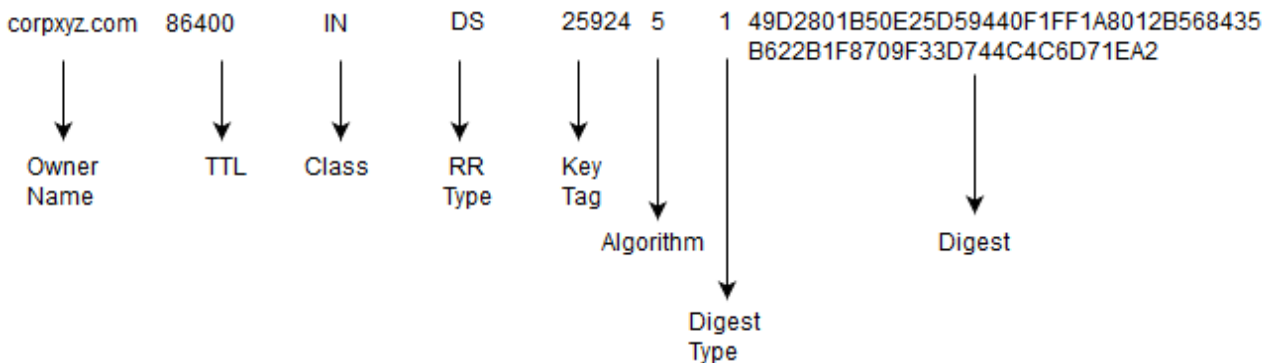
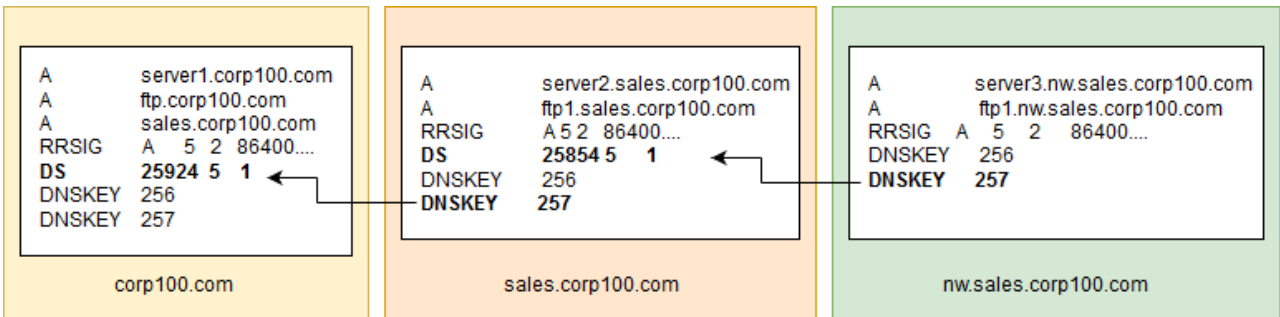
The first four fields specify the owner name, TTL, class and RR type. The succeeding fields are:

- Algorithm: The hash algorithm that was used. The currently supported algorithm is SHA-1, which is represented by a value of 1.
- Flags Field: Contains 8 one-bit flags, of which only one flag, the Opt-Out flag, is defined by RFC 5155. The Opt-Out flag indicates whether the NSEC3 record covers unsigned delegations.
- Iterations: The number of times the hash function was performed. The number of NSEC3 iterations is set to 10.
- Salt Field: A series of case-insensitive hexadecimal digits. It is appended to the original owner name as protection against pre-calculated dictionary attacks. New salt value is generated when the ZSK rolls over, for which the user can control the period. For random salt, the selected length is between one and 15 octets.

## DS Resource Records

A DS RR contains a hash of a child zone's KSK and can be used as a trust anchor in some security-aware resolvers and to create a secure delegation point for a signed subzone in DNS servers. As illustrated in the figure below, the DS RR in the parent zone corpxyz.com contains a hash of the KSK of the child zone sales.corpxyz.com, which in turn has a DS record that contains a hash of the KSK of its child zone, nw.sales.corpxyz.com.

Figure 22.1



The first four fields specify the owner name, TTL, class and RR type. The succeeding fields are as follows:

- **Key Tag:** The key tag value that is used to determine which key to use to verify signatures.
- **Algorithm:** Identifies the algorithm of the DNSKEY RR to which this DS RR refers. It uses the same algorithm values and types as the corresponding DNSKEY RR.
- **Digest Type:** Identifies the algorithm used to construct the digest. The supported algorithms are:
  - 1 = SHA-1
  - 2 = SHA-256
- **Digest:** If SHA-1 is the digest type, this field contains a 20 octet digest. If SHA-256 is the digest type, this field contains a 32 octet digest.

## Configuring DNSSEC on a Grid

You can configure the name servers in a Grid to support DNSSEC. You can configure the Grid Master as the primary server for a signed zone and the Grid members as secondary servers. (For more information, see [Configuring Grid Members to Support DNSSEC as Secondary Servers](#).) Note that only the Grid Master can serve as the primary server for a signed zone. Only the Grid Master can re-sign DNSSEC keys, and it must be primary for signed zones (as well as for signed zones with DNS Traffic Control records). That is, the Grid Master must serve the whole part of the configuration in which DNSSEC is used. Therefore, a running DNS service, DNS and DNS Traffic Control licenses are required on the Grid Master to use DNSSEC.

You can enable the Grid Master to sign zones and manage the DNSSEC keys, or you can configure the Grid Master as a client to a third-party, network-attached Hardware Security Module (HSM) that performs the key generation, zone signing, and key safekeeping. You must use either the Grid Master or HSM for zone signing and key management; you cannot use both. Note that each method may have different performance implications, depending on the hardware platform, number of zones and other factors. For information about using HSMs, see [About HSM Signing](#).

Any authoritative forward-mapping or reverse-mapping zone can be signed according to the following criteria:

- The zone does not contain any bulk host records.
- DNSSEC is enabled on the Grid Master.
- The primary server of the zone must be a Grid member. If the zone is assigned to an NS group, the primary server in the group must be a Grid member that has DNSSEC enabled.

Note that you can use DNS views to separate internal and external zone data, to manage your zones more efficiently and reduce the size of the zones that require signing. For information about DNS views, see [About DNS Views](#).

### Grid Master as Primary Server

When you sign a zone whose primary server is a Grid member, that member becomes a secondary server and the Grid Master becomes the hidden primary server. If the zone is assigned to an NS group, the Grid Master removes the association with the NS group. The previous primary server becomes a secondary server for the zone.

If a Master Candidate is promoted to Grid Master and the previous Grid Master was the primary server for signed zones, the new Grid Master becomes the hidden primary server for all signed zones. The previous Grid Master, which was the primary server for the zone, becomes a secondary server for the zone.

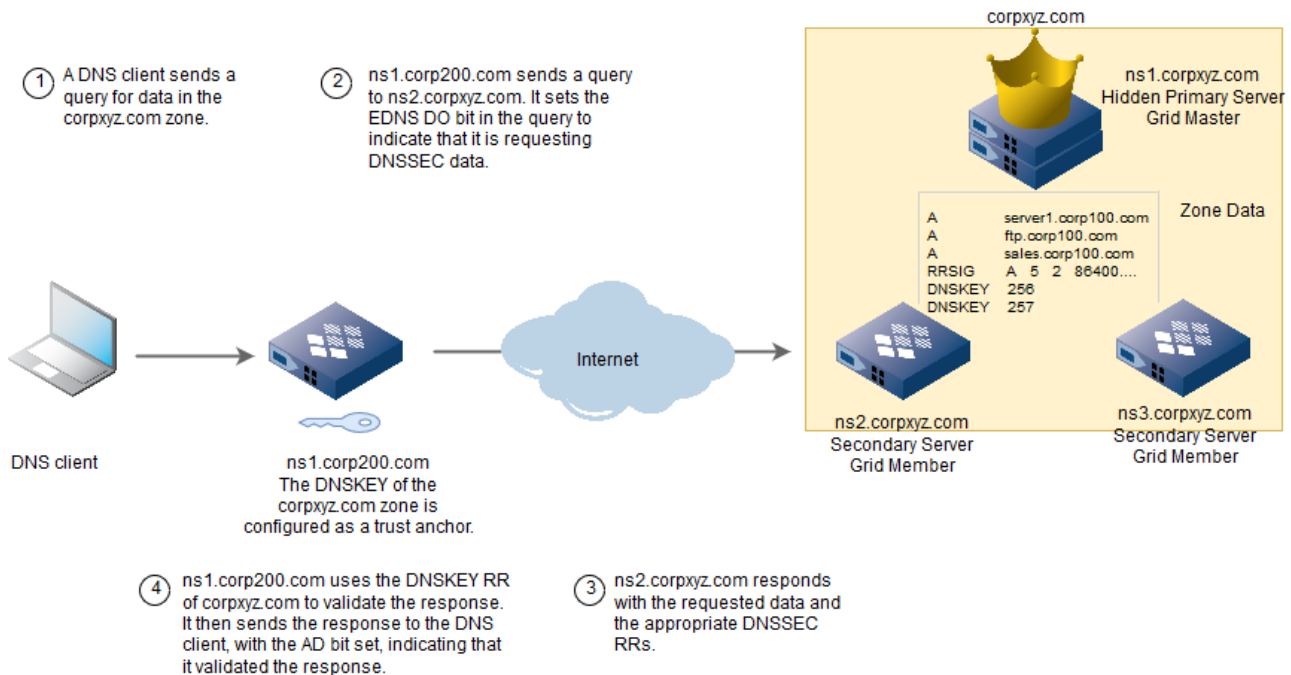
As the primary server, the Grid Master sends zone data to the secondary servers through zone transfers; or, if the secondary servers are Grid members, the Grid Master transfers data to all Grid members through the database replication process, by default. The Grid Master transfers all records in that zone, including all NSEC/NSEC3, RRSIG, DNSKEY and DS records with owner names that belong to that zone. The RRSIG RRs are included in zone transfers of the zone in which they are authoritative data. The Grid Master also performs incremental zone transfers to secondary servers as a result of incremental zone signings.

In addition, the Grid Master automatically performs an incremental signing of the zone data sets when their contents change. Incremental signing refers to signing just those parts of a zone that change when RRs are added, modified, or deleted. The Grid Master uses the private key of the ZSK when it incrementally signs a zone. In addition, the Grid Master adds, modifies or deletes the corresponding RRSIG records and the appropriate NSEC/NSEC3 records.

For example, the following Figure shows a Grid Master as the primary server of a signed zone and its Grid members as secondary servers. The Grid Master, ns1.corpxyz.com, is the hidden primary DNS server for the corpxyz.com zone. As the hidden primary name server for corpxyz.com, the Grid Master does not respond to queries from other name servers. Instead, it provides data to its secondary servers, ns2.corpxyz.com and ns3.corpxyz.com, which use this data to respond to DNS queries. Because the secondary servers are Grid members, they receive zone data from the Grid Master through the Grid database replication process.

The name server ns1.corp200.com is a recursive name server. It has configured the DNSKEY of the corpxyz.com zone

as a trust anchor. Therefore, it is able to validate the data it receives when it sends a query for the corpxyz.com zone.



Following are the tasks to configure the Grid Master to sign zones:

1. Create the zones. For information, see [Configuring Authoritative Zones](#).
  - Specify the Grid Master as the primary server.
2. Enable DNSSEC, as described in [Enabling DNSSEC](#).
3. Optionally, change the default DNSSEC settings. For information, see [Setting DNSSEC Parameters](#).
4. Sign the zone. The appliance automatically generates the DNSSEC RRs when you sign a zone. For information, see [Signing a Zone](#).

## Enabling DNSSEC

You can enable DNSSEC on a Grid, individual members, and DNS views. Because only Grid Masters can serve as primary servers for signed zones, you must enable DNSSEC on the Grid Master before you can sign zones. You must also enable DNSSEC on any Grid member that serves as a secondary server for signed zones.

When you enable DNSSEC on a Grid, you can set certain parameters that control the DNSSEC RRs, as described in [Setting DNSSEC Parameters](#).

When you enable DNSSEC on a Grid member or DNS view, you can set parameters that affect its operations as a secondary server, as described in [Configuring Grid Members to Support DNSSEC as Secondary Servers](#).

To enable DNSSEC on a Grid, member or DNS view:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **Members** tab -> *member* checkbox and click the Edit icon.  
**DNS View:** From the **Data Management** tab, select the **Zones** tab -> *dns\_view* checkbox and click the Edit icon.
2. In the editor, click **Toggle Expert Mode**.
3. When the additional tabs appear, click **DNSSEC**.
4. In the **DNSSEC** tab, select **Enable DNSSEC**.  
 Note when you disable EDNS0, all outgoing DNSSEC queries to zones within trusted anchors will fail even if DNSSEC validation is enabled. This is due to the restriction of the UDP packet length when you disable EDNS0. For information about EDNS0, see [Using Extension Mechanisms for DNS \(EDNS0\)](#).
5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Setting DNSSEC Parameters

The Grid Master uses certain default parameters when it signs a zone and generates the DNSSEC RRs. If you want to use different parameters for certain zones, you can change these defaults for the entire Grid or for individual zones. The following sections describe the different parameters that you can set:

- [About the DNSKEY Algorithm](#)
- [About Key Rollovers](#)
  - [Zone-Signing Key Rollover](#)
  - [Key-Signing Key Rollover](#)
  - [About Key Rollovers and DNS TTLs](#)
- [RRSIG Signatures](#)
- [Configuring DNSSEC Parameters](#)
- [Applying the Algorithm Changes](#)
- [Deleting the DNSSEC Keys Associated with a Zone](#)

For information on setting these parameters, see [Configuring DNSSEC Parameters](#) below.

### About the DNSKEY Algorithm

You can add multiple cryptographic algorithms that the Grid Master uses when it generates the KSK and ZSK. You can also add multiple algorithms at the Grid level and override them at the zone level. By default, the appliance uses RSA/SHA1 for both KSK and ZSK. You can add DSA, RSA/MD5, RSA/SHA1, RSA/SHA-256, RSA/SHA-512, ECDSAP/SHA-256, or ECDSAP/SHA-384 algorithms. However, you cannot add RSA/MD5 if the zone is signed with NSEC3 records. Note that a zone can contain either NSEC or NSEC3 records, but not both. You can add same or different set of algorithms with possibly different key sizes for both KSK and ZSK.

You can assign the DNSKey algorithm for HSMs. Entrust nShield HSMs don't support DSA. All other parameters are not used by HSMs.

### About Key Rollovers

To reduce the probability of being compromised, ZSKs and KSKs must be periodically changed. The time within which a key pair is effective is its rollover period. The rollover period starts as soon as a zone is signed. After a rollover period starts, you cannot interrupt or restart it unless you unsign the zone. During the key rollover, all the algorithms are rolled over at the same time and the rollover is performed on a per-zone basis.

### Zone-Signing Key Rollover

You can configure automatic ZSK rollovers on the Grid Master by using the double-signature rollover method or the pre-publish method. For more information, see [Configuring DNSSEC Parameters](#) below. The appliance initiates the ZSK rollover of signed zones when they are due. You can also perform a manual rollover of ZSKs. For more information about rolling zone-signing keys, see [Signing a Zone](#).

The double signature method provides a grace period, which is half of the rollover period. The default ZSK rollover period is 30 days; thus, the default grace period is 15 days.

At the end of a rollover period of a ZSK, the Grid Master generates a new ZSK key pair. It signs the zone with the private key of the new ZSK key pair, and consequently generates new RRSIG RRs with the new signatures. However, the Grid Master also retains the old ZSK key pair and RRSIG RRs. Thus, during the grace period, the data in the zone is signed by the private keys of both the old and new ZSKs. Their corresponding public keys (stored in DNSSEC RRs) can be used to verify both the old and new RRSIGs.

The grace period also allows the data that exists in remote caches to expire and during this time, the updated zone data can be propagated to all authoritative name servers. The Grid Master removes the old ZSK and its RRSIGs when the rollover grace period elapses. When a scheduled DNSSEC operation exists for a zone, the appliance does not lock it against other administrative changes and the administrator can still operate on a given zone even if there is a pending DNSSEC operation scheduled for it.

The appliance sets pre-publish method described in RFC 4641 as the default zone-signing key rollover method for NIOS 6.11.0 or later releases. In the pre-publish rollover method, the new key is published in the keyset before the actual



rollover. After the key propagates to all client caches, Grid Master removes the old signatures and creates new signatures with the new keys. The pre-publish rollover method uses the current key to sign the zone.

## Key-Signing Key Rollover

You can configure automatic KSK rollovers on the Grid Master or perform a manual KSK rollover. The default KSK rollover period is one year. The Grid Master also uses the double signature rollover method described in RFC 4641 for KSK rollovers. To configure automatic KSK rollovers, see as described in [Configuring Automatic KSK Rollovers and Notifications](#). For information about performing a manual KSK rollover, see [Rolling Zone-Signing Keys](#).

When the KSK rollover is overdue or is due within seven days, the Grid Master displays a warning when admins log in. In addition, you can also check which KSKs are due for a rollover as described in [Checking Key-Signing Keys](#).

When a user initiates a KSK rollover, the Grid Master sets the grace period to half the KSK rollover period. It generates a new KSK, and signs the DNSKEY records with the new KSK. Thus, during the grace period, the DNSKEY records are signed by the private keys of both the old and new KSKs. Both the old and the new KSKs can be used to validate the zone. The grace period allows the old keys in remote caches to expire. In addition, the admin should also export the new KSK and send it to the recursive name servers that use the KSK as trust anchors.

If the KSK rollover is for a child zone and the primary server of the parent zone is a Grid member, the Grid Master also inserts a DS record in the parent zone for the new DNSKEY in the child zone. If the primary server of the parent zone is external to the Grid, the admin must export either the DS record or the new KSK to the admin of the parent zone. For information about exporting a KSK, see as described in [Exporting Trust Anchors](#).

The Grid Master then removes the old KSK and its RRSIG records when the grace period for the KSK rollover ends.

## About Key Rollovers and DNS TTLs

Note that the KSK and ZSK rollover intervals affect TTLs used by RRs in signed zones.

A grace period is half of the key rollover interval. For example, if the KSK rollover interval is 1 year (365 days), then the grace period is 182.5 days; if the ZSK rollover interval is 30 days, then the grace period is 15 days.

The DNSKEY RRset in the zone is assigned a TTL, which is half of the signature validity interval. The default signature validity interval is set to 4 days, so DNSKEY RRset TTL is set to 2 days (172800 seconds).

All other RRs in the signed zone is limited to a "zone maximum TTL," which is the grace period of the ZSK. In the example, this is also 15 days.

When the zone is initially signed, if the TTL of an RR exceeds the zone maximum TTL, the Grid Master reduces the TTL to the zone maximum TTL. Additionally, the TTL settings for the signed zone are set to override; the values are inherited from the Grid DNS properties at that time, and the default TTL setting is reduced to the zone maximum TTL if the Grid property exceeds it. If the zone is later unsigned, the zone DNS properties remain at their overridden settings.

## RRSIG Signatures

As shown in the sample RRSIG record in [RRSIG Resource Records](#), the signatures have an inception and an expiration time. The default validity period of signatures in RRSIG records on the Grid Master is four days. You can change this default as long as it is not less than one day or more than 3660 days. The Grid Master automatically renews signatures before their expiration date.

## Configuring DNSSEC Parameters

The guidelines for choosing the DNSSEC parameters are as follows:

- RSA/SHA1 is the most widely used cryptographic algorithm for generating KSK and ZSK. However, it is recommended to use RSA/SHA-256 and RSA/SHA-512 for better interoperability.
- The usage of DSA cryptographic algorithm is optional. As stated in RFC 6944, it may not be supported by many systems.
- It is not recommended to use RSA/MD5 cryptographic algorithm as it is not very secure. As stated in RFC 6944, there are known defects in MD5.
- The key size of KSK algorithm is recommended to be equal to or greater than the key size of ZSK algorithm.

To set parameters at the Grid or zone level, complete the following steps:



1. **Grid:** On the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **Grid DNS Properties**.  
**Zone:** On the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox, and then click the Edit icon. Click **Override** to override the parameters.  
**Standalone appliance:** On the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **System DNS Properties**.
2. In the editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click **DNSSEC**.
4. On the **DNSSEC** tab, click the **Basic** tab and complete the following:
  - **Resource Record Type for Nonexistent Proof:** Select the resource record type (**NSEC** or **NSEC3**) you want to use for handling non-existent names in DNS. The default is **NSEC3**. The algorithms used by the KSK and ZSK can generate the same type of NSEC record. Note that a zone cannot contain both NSEC and NSEC3 resource records.
  - **Key-signing Key:** Click the Add icon to add the cryptographic algorithm that the Grid Master or HSM uses when it generates the KSK. You can add multiple algorithms, but you cannot add the same algorithm more than once. Grid Manager adds a row to the table each time you click the Add icon. Select the row and the algorithm from the drop-down list and enter the key size for the algorithm. The default is **RSA/SHA1** with the key size as **2048**.  
 Following are the valid values for each algorithm:  
**DSA:** The minimum is 512 bits and the maximum is 1024 bits, which is also the default. The key length must be a multiple of 64. Note that Entrust nShield HSMs do not support DSA.  
**RSA/MD5:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 2048 bits. You can configure this for NSEC only.  
**RSA/SHA1:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 2048 bits.  
**RSA/SHA-256:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 2048 bits.  
**RSA/SHA-512:** The minimum is 1024 bits, the maximum is 4096 bits, and the default is 2048 bits.  
**ECDSAP/SHA-256:** The minimum is 160 bits, the maximum is 256 bits.  
**ECDSAP/SHA-384:** The minimum is 160 bits, the maximum is 384 bits.  
 You can delete an algorithm by selecting it and clicking the Delete icon.
  - **Key-signing Key Rollover Interval:** Specify the key signing key rollover interval for all the algorithms. The minimum value is one day and the maximum is the time remaining to January 2038. The default is one year.
  - **Zone-signing Key:** Click the Add icon to add the cryptographic algorithm that the Grid Master or HSM uses when it generates the ZSK. You can add multiple algorithms, but you cannot add the same algorithm more than once. Grid Manager adds a row to the table each time you click the Add icon. Select the row and the algorithm from the drop-down list and enter the key size for the algorithm. The default is **RSA/SHA1** with the key size **1024**.  
 Following are the valid values for each algorithm:  
**DSA:** The key length must be a multiple of 64. The minimum is 512 bits and the maximum is 1024 bits. The default is 1024 bits. Note that Entrust nShield HSMs do not support DSA.  
**RSA/MD5:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 1024 bits. You can configure this for NSEC only.  
**RSA/SHA1:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 1024 bits.  
**RSA/SHA-256:** The minimum is 512 bits, the maximum is 4096 bits, and the default is 1024 bits.  
**RSA/SHA-512:** The minimum is 1024 bits, the maximum is 4096 bits, and the default is 1024 bits.  
**ECDSAP/SHA-256:** The minimum is 160 bits, the maximum is 256 bits.  
**ECDSAP/SHA-384:** The minimum is 160 bits, the maximum is 384 bits.  
 You can delete an algorithm by selecting it and clicking the Delete icon.
  - **Zone-signing Key Rollover Interval:** Specify the zone signing key rollover interval for all the algorithms. The minimum value is one day and the maximum is the time remaining to January 2038. The default is 30 days.
  - **Signature Validity:** Specify the signature validity period for RRSIG RRs. The minimum is one day and the maximum is 3660 days. The default signature validity interval is four days.
  - **Zone-signing Key rollover method:** You can use either of these methods to sign all the RRsets in a zone:
    - i. **Pre-publish:** Select this if you want to use the pre-publish signature scheme to sign all the RRsets in a zone while performing the ZSK rollover. When you select this option, the record sets are signed using a single key. The appliance sets this option as the default zone-signing key method for NIOS 6.11.0 and later releases.

- ii. **Double Sign:** Select this if you want to use the double signature scheme to sign all the RRsets in a zone while performing the ZSK rollover. The non-DNSKEY RRset are signed twice, which increases the size of the zone files.

Note that you can select the **Zone-signing Key rollover method** only after you enable DNSSEC.

5. Save the configuration and click **Restart** if it appears at the top of the screen.

When you modify the algorithms for a signed zone, you can apply the algorithm changes to the zone, as described later or you can unsign the zone and sign it again. For an unsigned zone however, you can apply the algorithm changes by signing the zone. For information about signing a zone, see [Signing a Zone](#).

When you re-sign a zone after adding an algorithm, the DNSKEY key pairs of the old algorithms are rolled over and all the old RRSIG records are removed. The zone is re-signed with the new DNSKEY key pairs. When you re-sign a zone after removing an algorithm, the DNSKEY key pairs of the remaining algorithms are rolled over and the DNSKEY key pairs of the removed algorithm is removed. All old RRSIG records are removed and the zone is re-signed with the new DNSKEY key pairs.



#### Note

If you add or remove a KSK algorithm from a zone, you must update the DS RRsets at the parent zone when the parent zone is managed by a non-Infoblox DNS server or by an Infoblox server that is part of a different Grid. For information, see [Importing a Keyset](#).

## Applying the Algorithm Changes

You can apply the algorithm changes to a zone whenever the KSK or ZSK algorithms are modified. You can apply the algorithm changes only to a signed zone.

To apply the algorithm changes to a signed zone, complete the following steps:

1. On the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Apply Algorithm Changes**.
3. In the *Apply Algorithm Changes* dialog box, click the Add icon to select a zone. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. The appliance displays signed zones only. Select a zone. To add multiple zones, click the Add icon and select a zone.  
To remove a zone from the list, select the checkbox adjacent to the respective zone and click the Delete icon. You can click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, either select **Now** or select **Later** and enter a date, time, and time zone.
4. After you have selected the zones, click **Apply Algorithm Changes**.
5. When the confirmation dialog appears, click **Yes**.

## Deleting the DNSSEC Keys Associated with a Zone

You can view the status of KSKs and ZSKs or delete the existing keys. Note that you can only delete keys that are either published or rolled over. You cannot delete keys that are active.

To delete keys, complete the following steps:

1. On the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox, and click the Edit icon.
2. In the *Authoritative Zone* editor, click the **DNSSEC** tab, and the following fields are displayed in the **Advanced** tab:
  - **Key ID:** The appliance generates a separate ID for each key. Select the checkbox adjacent to the **Key ID**. To delete a KSK, select the checkbox in the **Key-signing Keys** table. Select the checkbox in the **Zone-signing Keys** table to delete a ZSK.
  - **Status:** The **Status** column displays the status of the respective key. It can be one of the following: **Active**, **Published**, or **Rolled**.
  - **Public Key:** This column displays the public key that is associated with the respective KSK or ZSK.
  - **Algorithm:** This column displays the algorithm that is associated with the respective KSK or ZSK.
  - **Time until next event:** This column displays the time that is left to perform the next action for a key that is associated with the respective ZSK. This column helps you decide whether to roll over manually or wait for a zone to resign automatically. The time is displayed in months, days, hours format. For example, 2m 3d 13h implies time left to perform the next action is 2 months, 3 days and 13 hours.
    - **Active Key:** Indicates the time when the active key is rolled and zone is signed with the published key.

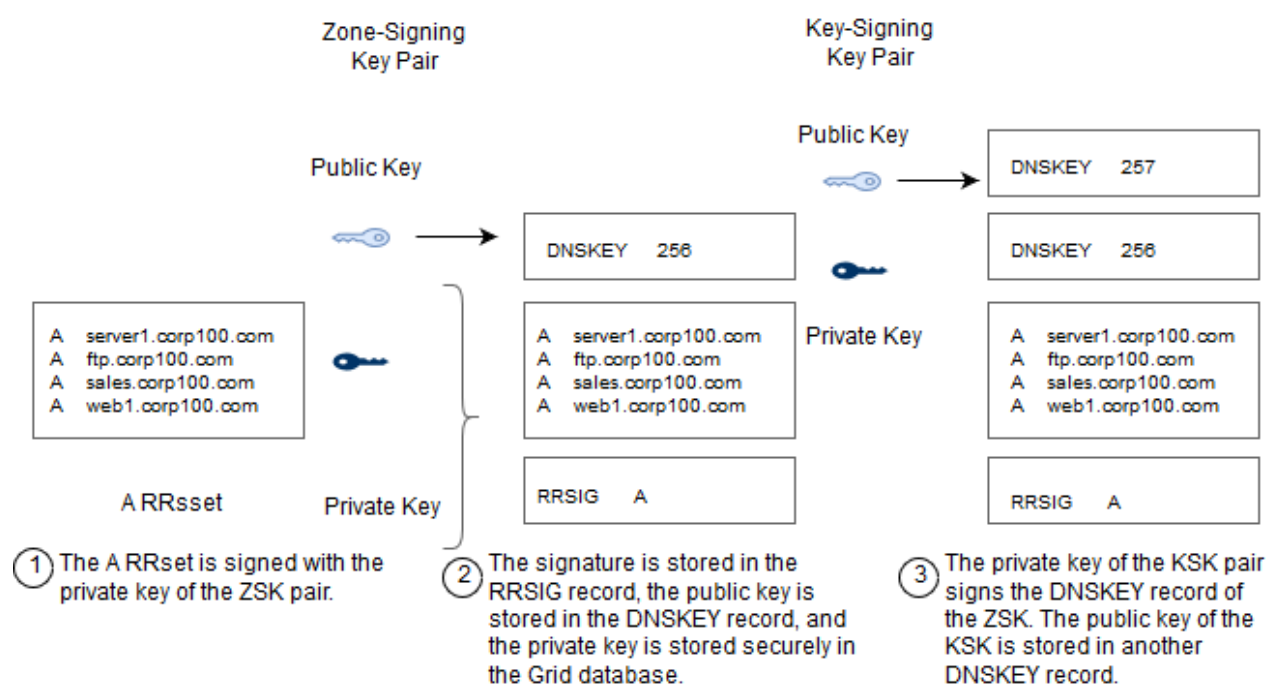
- **Published Key:** Indicates the time when the published key is used to resign a zone.
- **Rolled Key:** Indicates the time when a rolled key is deleted. Rolled keys are stored for quite a long period of time and are not used. You can manually cleanup the rolled keys.

3. Click the Delete icon.

## Signing a Zone

When it signs a zone, the Grid Master generates new DNSKEY key pairs. As shown in the below figure, it uses the private key of the ZSK to sign the authoritative RRsets in the zone, and stores the corresponding public key in a DNSKEY record. It then uses the private key of the KSK to sign the DNSKEY records and stores the corresponding public key in another DNSKEY record. It stores the private keys in the Grid database and stores the public keys in the DNSKEY records in the database.

### Zone Signing Process



The Grid Master also does the following:

- It inserts NSEC or NSEC3 records. The use of NSEC or NSEC3 records depends on the NSEC type you selected for the Grid or the zone. When you select NSEC3, the Grid Master uses NSEC3 records in signed zones.
- It increments the SOA serial number and notifies the secondary servers that there is a change to its zone data. When the secondary servers check the serial number and see that it has been incremented, the secondary servers request a zone transfer.
- If the TTL of an RR in the zone exceeds the ZSK grace period, the Grid Master reduces the TTL to the ZSK grace period. (For information about the grace period, see [About Key Rollovers](#).) Setting a TTL value that exceeds half of the rollover period is not allowed.
- If the KSK rollover period is less than the ZSK rollover period, the Grid Master sets the TTL of the DNSKEY RR to the KSK rollover period.
- The appliance sets the Grid Master as the primary server for zones, enables DNSSEC on the Grid Master, and starts DNS service on the Grid Master.

When it signs a subzone, the Grid Master automatically inserts DS records for parent zones that are hosted by Grid members. The appliance allows you to sign a single zone or multiple zones simultaneously. For example, if you have multiple zones that are due for rollover at the same time, you can select all such zones and sign them at once. Note that each operation is independent of the other. For example, if you want to sign five zones at the same time, and if one of the zones fails during this time, NIOS signs the remaining four zones. Note that the selected zones must have an associated

primary server. The appliance displays an error message if the zone does not have a primary server. When the sign operation fails, the appliance displays the zone names, associated DNS views, and the error message indicating the reason for failure.

To sign a zone:

1. From the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Sign Zones**.
3. In the *Sign Zones* dialog box, the displayed zone name can either be the last selected zone or the zone from which you are signing. If no zone name is displayed or if you want to select a different zone, click the Add icon. The appliance displays unsigned zones only. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Select a zone. To add multiple zones, click the Add icon and select a zone. You can click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, either select **Now** or select **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#).
4. After you have selected the zones, click **Sign Zones**.
5. When the confirmation dialog displays, click **Yes**.

When you sign multiple zones, the appliance displays generic error messages for the following cases:

- The value to which the resource record TTL is reduced is not displayed.
- The appliance displays a message about name server group disassociation if at least one zone is associated with a name server group. It will not list the affected zones.
- When you sign a zone or multiple zones, the appliance displays a warning message indicating that the operation might take a longer time.
- The appliance displays an error message if the number of characters in the zone name, which you want to sign, exceeds 180 characters. You can sign a zone only when the name of the zone is less than 180 characters in size.

To remove a zone from the list, select the checkbox adjacent to the respective zone and click the Delete icon. To view the records of the signed zone, from the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone*. Expand the Records section to list the RRs of the zone, as shown in the below figure.

The screenshot shows the 'Records' section of a DNS management interface. It features a toolbar with icons for adding, deleting, and refreshing records, along with a 'Go to' search field. Below the toolbar is a table with columns: NAME, TYPE, DATA, RECORD SOURCE, PRINCIPAL, PROTECTED, COMMENT, and MONITORED SINC. The table contains several rows of RRSIG and NSEC3PARAM records, all with a 'System' record source and 'Not Monitored' status.

NAME	TYPE	DATA	RECORD SOURCE	PRINCIPAL	PROTECTED	COMMENT	MONITORED SINC
	RRSIG Record	NS 8 2 28800 20181214080354 2018...	System				Not Monitored
	RRSIG Record	DNSKEY 8 2 172800 201812140803...	System				Not Monitored
	RRSIG Record	DNSKEY 8 2 172800 201812140803...	System				Not Monitored
	RRSIG Record	NSEC3PARAM 8 2 900 20181214080...	System				Not Monitored
	NSEC3PARAM R...	1 0 10 2185CA421135	System				Not Monitored
	DNSKEY Record	257 3 8 AwEAAe1cf5wu56JkaaAq1...	System				Not Monitored
	DNSKEY Record	256 3 8 AwEAAZ+8b+qYIDYeEb3V...	System				Not Monitored

## Managing Signed Zones

After you sign a zone, you can do the following:

- You can add a DS RR at the delegation point for a signed subzone when the subzone is hosted on a non-Infoblox DNS server or an Infoblox server that is part of a different Grid. For information, see [Importing a Keyset](#) below.
- Trust anchors can be specified as DNSKEY RRs, DS RRs, and as a BIND trusted-keys statement. You can export any of these as trust anchors. For information, see [Exporting Trust Anchors](#) below.
- You must change the KSK periodically, to ensure its security. For information, see [Checking Key-Signing Keys and Rolling Key-Signing Keys](#) below.
- You can initiate ZSK rollovers manually. For information, see [Rolling Zone-Signing Keys](#) below.

- If, for any reason, the security of the keys are compromised, you can delete a key and perform a manual rollover. For information, see [Configuring Emergency KSK Rollover](#) below. Note that when you re-sign a zone, the Grid Master generates new ZSK and KSK pairs. You must send the new DNSKEY of the KSK to resolvers that use it as a trust anchor and generate new DS records and send them to the parent zones.
- You can move a signed zone to the Recycle Bin, from where you can delete it permanently or restore it. For information, see [Deleting and Restoring Signed Zones](#) below.

In addition, signed zones can accept dynamic DNS updates. For information about configuring zones to accept dynamic DNS updates, see [Configuring DNS Servers for DDNS](#).

### Importing a Keyset

A keyset is a DS RRset, or a DNSKEY RRset which is used as input to generate the DS RRset. To import a keyset:

1. From the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Import Keyset**.
3. In the *Import Keyset* dialog box, the displayed zone name can either be the last selected zone or the zone from which you are importing the keyset. If no zone name is displayed or if you want to select a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box from which you can select a zone.
4. Paste the KSK or DS record being imported. It must be a KSK or DS record, and must belong to an immediate subzone of the zone to which the record is being imported.
5. Click **Import**.

If you imported a DNSKEY RRset, the Grid Master uses the SHA-1 algorithm to generate the DS RRset, which it adds to the parent zone. If you imported a DS RRset, the Grid Master adds it to the parent zone. The Grid Master incrementally signs the DS RRset.

### Exporting Trust Anchors

A trust anchor is a DNSSEC public key which is used by security-aware resolvers as the starting point for establishing authentication chains. A trust anchor can be specified as a DNSKEY RR or a DS RR, which contains the hash of a DNSKEY RR and can also be used to create a secure delegation point for a signed subzone in DNS servers. In BIND, trust anchors are configured using the trusted-keys directive. A trusted key is a DNSKEY RR without the TTL, class and RR type. You can export the trust anchors for the selected zone in a format that can be used in a BIND trusted-keys directive. Exporting trust anchors supports multiple algorithms, which means you can export all the algorithms in a key.

To export trust anchors:

1. From the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Export Trust Anchors**.
3. In the *Export Trust Anchors* dialog box, do the following:
  - The displayed zone name can either be the last selected zone or the zone from which you are exporting trust anchors. If no zone name is displayed or if you want to select a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box from which you can select one.
  - Select one of the following: **DNSKEY records**, **DS records**, or **BIND trusted-keys statement**.
4. Click **Export**.
5. Specify the location of the exported file and click **OK**.

If you exported DS records, the exported file contains DS records that use the SHA-1 and SHA-256 algorithms.

### Checking Key-Signing Keys

To check which key-signing keys are overdue for a rollover or are due to roll over within a week:

1. From the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Check KSK Rollover Due**.

3. The *KSK Rollover Due* dialog box lists the key-signing keys that are due to rollover. It includes the domain name of the zone, DNS view (if there are multiple DNS views), and the number of days until the rollover.
4. You can click the Schedule icon at the top of the wizard to schedule a KSK rollover for one or more zones at a given date and time. In the *Schedule Change* panel, either select **Now** or select **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#).
5. Click **Close**.

## Rolling Key-Signing Keys

You can initiate a rollover before or after a rollover period, or when you need to replace the KSK for security reasons. You can initiate a KSK rollover several times simultaneously, but note that the number of keys will increase each time you perform a rollover. You can schedule the KSK rollover to occur at a later date and time.

To roll over key-signing keys:

1. From the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Roll Over Key-Signing Key**.
3. In the *Roll Over Key-Signing Key* dialog box, the displayed zone name can either be the last selected zone or the zone from which you are rolling over key-signing keys. If no zone name is displayed or if you want to select a different zone, click the Add icon. The appliance displays signed zones only. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. To add multiple zones, click the Add icon and select a zone. You can click the Schedule icon at the top of the wizard to schedule a KSK rollover for one or more zones at a given date and time. In the *Schedule Change* panel, either select **Now** or select **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#). Note that you cannot schedule the KSK rollover on a recurring basis.
4. Click **Roll Over**.

You can export the new KSK and send it to the security-aware resolvers that use it as a trust anchor. To remove a zone from the list, select the checkbox adjacent to the respective zone and click the Delete icon.

## Rolling Zone-Signing Keys

Only an administrator can initiate ZSK rollovers either before or after a rollover period, or when you want to replace the ZSK for security reasons. You can initiate a ZSK rollover several times simultaneously, but note that the number of keys will increase each time you perform a rollover.

To roll over zone-signing keys:

1. From the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Roll Over Zone-Signing Key**.
3. In the *Roll Over Zone-Signing Key* dialog box, the displayed zone name can either be the last selected zone or the zone from which you are rolling over zone-signing keys. If no zone name is displayed or if you want to select a different zone, click the Add icon. The appliance displays unsigned zones only. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. To add multiple zones, click the Add icon and select a zone. You can click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, either select **Now** or select **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#).
4. Click **Roll Over**.

To remove a zone from the list, select the checkbox adjacent to the respective zone and click the Delete icon. The appliance displays warning messages when the changes take effect on the first zone or when the rollover occurs. You cannot change the zone-signing key rollover method while the previous change is still in progress. The previous change will be completed only when the zone active key, which is active when the rollover method is changing, expires and is deleted.

## Best Practices for Configuring Zone Signing Keys

Infoblox recommends that you use the pre-publish option for zone signing key method for the following reasons:

- The double-signature ZSK rollover doubles the number of signatures in your zone when a rollover is in progress. The size of the zone increases due to the duplicate signature records. This is not recommended if the size of your

zones are large. When you select this option, the appliance creates a new set of signatures for all the resource records. This also increases the database usage.

- When you select to pre-publish key rollover, the rollover uses a single key to sign the records at a given time and it does not sign the zone data twice. The appliance publishes the new key in the keyset even before the actual rollover. This reduces the database usage.

## Unsigning a Zone

When you unsign a zone, the Grid Master permanently removes all automatically generated DNSSEC records in the zone and parent zone. It does not remove any DS records associated with a child zone. You can unsign a single zone or multiple zones at the same time.

To unsign a zone:

1. From the **Data Management** tab, select the **DNS** tab.
2. Expand the Toolbar and click **DNSSEC -> Unsign Zones**.
3. In the *Unsign Zones* dialog box, the displayed zone name can either be the last selected zone or the zone from which you are unsigning. If no zone name is displayed or if you want to select a different zone, click the Add icon. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. The appliance displays signed zones only. Select a zone. To add multiple zones, click the Add icon and select a zone. You can click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, either select **Now** and click **Save** or select **Later** and enter a date, time, and time zone. For information, see [Scheduling Tasks](#).
4. After you have selected the zones, click **Unsign Zones**.
5. When the confirmation dialog displays, click **Yes**.

To remove a zone from the list, select the checkbox adjacent to the respective zone and click the Delete icon.

## Deleting and Restoring Signed Zones

When you delete a signed zone, the Grid Master unsigns the zone before moving it to the Recycle Bin. Unsigning the zone effectively deletes all auto-generated DNSSEC RRs; only user-defined DS records are retained and moved to the Recycle Bin as well. The Grid Master also retains the ZSK and KSK in its database, until you permanently delete the zone from the Recycle Bin.

When you restore a signed zone, the Grid Master restores it and re-signs its data sets with the original keys, which are also restored. You can also restore the user-defined DS records. The rollover period of the ZSK and KSK starts when the zone is signed after it is restored. Note that when you restore a zone that contains rolled keys, either KSK or ZSK, the appliance removes all these rolled keys.

Note that when you restore a deleted zone from recycle bin on the NIOS server, which is created and signed on the Microsoft Server 2012, then all the DNSSEC records will be deleted, except for the DNSKEY records. The DNSKEY records will only be resynchronized. The DNSSEC records are read-only and cannot be regenerated using NIOS. You must recreate the zone manually on the Microsoft Server. When you recreate the zone on the Microsoft Server, new keys will be generated. The signed zone, which is restored, and the DNSSEC keys are synced to Microsoft Server. This zone will be seen as an unsigned zone on the Microsoft Server, as NIOS does not trigger the signing zone request for the corresponding zone. For such zones, the 'DNSSEC' label is not displayed and the value for 'Signed' column is 'No'.

To delete a signed zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab.
2. Click the checkbox of the zone you want to delete.
3. Click the Delete icon.
4. Click **Yes** to confirm the deletion.

To restore a signed zone:

1. In the *Finder* panel, expand **Recycle Bin**.
2. Select the zone you want to restore.
3. Click the Restore icon.



## Configuring Automatic KSK Rollovers and Notifications

You can configure automatic KSK rollovers at the Grid level and override the settings at the zone level. You can also configure notifications for KSK rollovers. The appliance sends one notification, SNMP, or email, or both based on the selection, for each event. For example, if the KSK of two zones are rolled over in the same batch, the appliance sends two notifications, one for each zone. Note that the appliance sends these notifications only once, and they are not recurring. Apart from the notifications that you receive, Grid Manager also displays a banner when you log in to the Grid indicating that the KSK rollover is due within the next seven days.

These notifications are not applicable to an ZSK, as the ZSK rollover is an automated process. The appliance generates numerous notifications.

To configure KSK rollover, complete the following:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **Grid DNS Properties**.  
**Zone:** From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox, and click the Edit icon. Click **Override** to override the values at the zone level.  
**Standalone appliance:** From the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **System DNS Properties**.
2. Select the **DNSSEC** tab and complete the following in the **Basic** tab:
  - **KSK Notification Configuration:** You can choose to receive notifications for KSK rollover events.
    - **No Notifications:** Select this if you do not want to receive any notifications for KSK rollover events.
    - **Notifications for all KSK rollover events:** Select this if you want to receive notifications for all KSK rollover events. The appliance sends notifications after the rollover.
    - **Notifications only for KSK rollover events requiring manual DS update to parent zone:** Select this if you want to receive notifications only for KSK rollover events that require manual DS updates to parent zone. This is selected by default.
    - **Enable KSK Email Notification:** Select this to receive email notifications about DNSSEC keys.
    - **Enable KSK SNMP Notification:** Select this to receive SNMP trap alerts about DNSSEC keys.
3. **Enable automatic KSK rollover:** This is selected by default. When you select this option, the appliance will automatically roll over KSKs when they are due. The appliance starts the rollover process at most six hours after the due date. The appliance logs the messages in the syslog.  
Note that the appliance enables notifications and automatic KSK rollover by default for NIOS 6.11.0 and later releases.  
These are not available for earlier releases. Similar to automatic ZSK rollover, the appliance automatically restarts the DNS service after a KSK is rolled over.
4. Save the configuration.

## Configuring NSEC3 Salt Length and Hashing Iterations

The salt is a random string, which is appended to the domain name before it gets hashed. The number of iterations indicates the number of additional times the hashing occurs. These serve as a protection against dictionary attacks. The appliance generates a new salt for initial signing and changes it every time a ZSK rollover occurs. Note that when you use a longer salt and higher number of iterations, DNS is more secure and the chances of dictionary attacks on NSEC3 are reduced.

You can choose the minimum and the maximum salt length at the Grid level and override them at the zone level. Note that the length of the salt has an impact on the size of the NSEC3 record, but it does not have an impact on the performance of the appliance.

When the number of iterations increases, the DNS client has to validate a additional data and the cost of the DNS server to serve the zone increases. This might also reduce the performance of the system with regards to DNSSEC operations.

To define salt length and hashing iterations, complete the following:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **Grid DNS Properties**.  
**Zone:** From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox, and click the Edit icon. Click **Override** to override the parameters.  
**Standalone appliance:** From the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **System DNS Properties**.
2. Select the **DNSSEC** tab and complete the following in the **Basic** tab:
  - **Zone-signing Key Settings**



- **NSEC3 Salt Length:** Specify a minimum and maximum length for NSEC3 salt. The minimum length is one octet and the maximum length is 255 octets. The appliance sets the following default values for minimum and maximum lengths respectively: one and 15 octets.
- **Number of NSEC3 hashing iterations:** The appliance uses the default value, ten, for hashing iterations. The minimum value is ten and the maximum value depends on the smallest key size, as defined in RFC 5155 as follows:
  - 150 if the key size is equal or less than 1024 bits.
  - 500 if the key size is equal or less than 2048 bits.
  - 2,500 if the key size is equal or less than 4096 bits.



**Note**

The above fields are displayed only when you select NSEC3 record type.

## Deleting Server Keys

The appliance retains the key until the expiration of the grace period. For example, if the validity period of a KSK is two years, you can delete the rolled key after publishing the DS record to the parent zone and waiting for a period greater than its TTL.

The following rules are valid for KSK and ZSK signing using the double-signature scheme:

- You cannot delete an active key.
- When you delete a rolled key, the appliance displays a warning message indicating that it might break validation on clients.

The following rules are valid for ZSK signing using the pre-publish scheme:

- You cannot delete an active key.
- When you delete a pre-published key, the appliance generates a new pre-published key.
- You can delete a rolled key. The appliance deletes this key as it is no longer used.

When you use an HSM, the appliance does not delete the key from HSM. For more information, see [About HSM Signing](#).

## Configuring Emergency KSK Rollover

The appliance supports emergency rollover that can be used when the keys are compromised. In an emergency operation, you must delete the compromised key and the associated compromised data from the zone. The ability to perform emergency rollovers enable administrators to react quickly when a zone is compromised. To initiate an emergency rollover, you must first perform a manual rollover. For information about rolling over a KSK manually, see [Rolling Key-Signing Keys](#) below. After the rollover, you must delete the compromised key. For information about deleting the compromised key, see [Deleting Server Keys](#) below.

An emergency KSK rollover involves the following:

- The administrator of the compromised zone, which is hosted on the Infoblox appliance, must initiate the emergency KSK rollover and later export the corresponding DS record.
- The administrator of the parent zone, which is hosted on an external server, must import the DS record of the child zone. This is required to maintain the chain of trust.

During this emergency procedure, the chain of trust is temporarily broken. As stated in RFC 6781, the effect depends on the order of the operations:

- You must perform the KSK rollover first. The chain of trust is broken until the administrator of the parent zone replaces the DS record. In the meantime, the zone appears bogus to a validating client.
- You must remove the compromised DS record first. The chain of trust is broken until the NIOS administrator performs the KSK rollover, communicates the new DS record to the administrator of the parent zone who then adds it. In the meantime, the zone appears insecure to validating resolvers.

## Handling Error and Warning Messages for DNSSEC Operations

Grid Manager consolidates warnings from one or more DNSSEC long running operations and displays them in a single dialog box. For more information about long running tasks, see [About Long Running Tasks](#). When the long running operations are completed with warnings, the appliance displays the message using a yellow banner at the top of Grid Manager. Click **Show Warnings** to view the warnings that are generated by the server for DNSSEC operations. Click **Close** to close the warning dialog box.

The appliance also displays a warning icon in the **Execution Status** column of the **Task Manager**. For more information about the task manager, see [Viewing Tasks](#). When you click the icon, the appliance displays a dialog box indicating that the task is complete with warnings. Click the **show task details** hyperlink to view detailed information about the error or warning message. The appliance displays *Scheduled Task Details* wizard in which you can view the task details and error or warning messages, if any. For more information, see [Viewing Scheduled Tasks](#). The appliance clears the warning messages from the cache when you close the dialog box.

When a DNSSEC operation fails, NIOS saves the messages in the syslog. For more information about the syslog, see [Using a Syslog Server](#).

## Viewing Scheduled Tasks

You can view the status of operations that are scheduled. The appliance also displays error or warning messages, if any. For example, if you schedule the **Sign Zones** operation for multiple zones, there is a possibility that some operations may fail, some may succeed with warnings, and some are completed successfully. You can also view the error or warning messages that are generated for certain affected zones. To view the status:

1. From the **Administration** tab, select the **Workflow** tab -> **Task Manager** tab.
2. Grid Manager displays the DNSSEC operations that are scheduled. To view the details you can either click the Action icon  
  
next to the task ID and select **View** from the menu or select the checkbox adjacent to the Action icon and then click **View** from the Toolbar.
3. The **General** tab of the *Scheduled Task Details* wizard displays the following details:
  - **Task ID:** The ID associated with the task. The appliance assigns an ID to a task in chronological order. By default, the appliance sorts tasks by **Task ID**.
  - **Action Type:** The operation the appliance performs in this task.
  - **Submitter:** The username of the admin who scheduled or submitted the task.
  - **Submitted Time:** The date, time, and time zone when the task was submitted.
  - **Ticket Number:** For an approval workflow, this number may be entered by the submitter to associate the task with a help desk ticket number or a reference number.
  - **Approver:** The username of the admin who has approved this task.
  - **Approver Comments:** Comments entered by the approver.
  - **Executed on Member:** The Grid member on which the task is executed.
  - **Execution Status:** The execution status of the task. Possible values are **Completed**, **Failed**, **Pending**, and **Executing**.
  - **Execution Time:** The date, time, and time zone when the task was executed.
  - **Affected Objects:** The name of the object and object type.
4. The **Warnings/Errors** tab of the *Scheduled Task Details* wizard displays error or warning messages related to tasks. It also displays object execution details. This table is blank if there are no error messages or warnings. You can view the error message and the name of the zone with which the error or warning message is associated.
5. Click **Close** to close the dialog box.

## About HSM Signing

You can integrate a Grid with third-party, network-attached Hardware Security Modules (HSMs) for secure private key storage and generation, and zone-signing off-loading. Infoblox appliances support integration with either Thales Luna HSMs or Entrust nShield HSMs. When using a network-attached HSM, you can provide tight physical access control, allowing only selected security personnel to physically access the HSM that stores the DNSSEC keys. When you enable this feature, the HSM performs DNSSEC zone signing, key generation, and key safe keeping.

Note that if you migrate from using the Grid Master to HSMs, HSM signing starts at the next key rollover. Only a superuser can configure this feature. To configure HSM signing in a Grid, do the following:

1. Create the HSM group and add HSMs to the group. You can create either a Thales Luna HSM group or an Entrust nShield HSM group. You can use only one group at a time. After you add the HSM group, the Add icon and Add button in the Toolbar are greyed out.
  - For information on adding a Thales Luna HSM group, see [Configuring a Thales Luna HSM Device](#) below.
  - For information on adding an Entrust nShield HSM group, see [Adding and Managing an Entrust nShield HSM Group](#) below.

Note that if you delete an HSM or an HSM group, it is permanently deleted. It is not stored in the Recycle Bin.

2. Enable HSM signing. For information, see [Enabling HSM Signing](#) below.

After you enable this feature, you can monitor the HSM group, as described in [Monitoring the HSM Group](#) below. In addition, the Grid sends SNMP traps when zone signing succeeds or fails. For information about these traps, see [Processing and Software Failure Traps](#).

Note that NIOS does not provide key life cycle management functions; these are handled by the HSM and must be configured via the HSM's administrative interface to adhere to corporate policies on key management. The keys (ZSK and KSK) used for DNSSEC are stored securely on the HSM and are not deleted by NIOS when the key is no longer required by the DNSSEC function. However, references to the keys are removed from the appliance.

## Configuring a Thales Luna HSM Device

You can integrate a Grid with a Thales Luna HSM group. The Thales Luna HSM group can contain either Thales Luna 4, Luna 5, or Luna 6 devices in standalone or HA mode; the group cannot contain a mix of both models. You must first configure each HSM device, and then create the group and add the devices to the group, as described in [Adding a Thales Luna HSM Group](#) below.

### Configuring a Thales Luna HSM Device

Do the following for each Thales Luna HSM device that you are adding to the group:

1. On the Grid, generate a client certificate for the Grid Master and Grid Master Candidate. For information, see [About Client Certificates](#). If you are upgrading the Thales Luna HSM from Luna 5 or 6 to Luna 7 CPL, you must generate a new client certificate.
2. On the Thales Luna HSM, do the following:
  - Assign the Grid Master and Grid Master Candidate to a partition on the HSM to avoid any service interruptions, in case the Grid Master Candidate is promoted to Grid Master.
  - Upload the certificates of the Grid Master and Grid Master Candidate to the HSM and register the certificates in the HSM's list of clients. The certificates of the Grid Master and Grid Master Candidate are linked to their IP addresses. Therefore, if any of their IP addresses change, you must generate a new client certificate and register it with the HSM.

Note that if the HSM is configured and you replace an appliance that was a Grid Master or Grid Master Candidate and you backed up the database of the old appliance and restored it on the replacement appliance, the certificates remain intact. Therefore, you do not need to regenerate a new certificate for the replacement, as long as the IP address does not change.
  - If you are upgrading from a previous version of Thales Luna HSM to a later version, such as from Luna 6 to Luna 7 CPL, you must complete the following before adding the new Luna configuration to NIOS:
    - Remove the previous certificate registration from the HSM server and then re-register the Grid Master and Grid Master Candidate certificates.
    - Generate a new HSM certificate if you want to retain the current IP settings for the Grid Master.
  - Download the HSM certificate.



#### Note

- Make sure that the common name used in the certificates is distinct when you configure HSM servers in HA mode.
- To configure Thales Luna on an HA pair, add a static route with the virtual IP address of the Grid to the HSM server.

For additional information, refer to your Thales Luna HSM documentation.

## Adding a Thales Luna HSM Group

When you configure a Thales Luna HSM group, add the Thales Luna HSM devices to the group and upload their certificates to the Grid. You can add only one HSM group. To add a Thales Luna HSM Group:

1. From the **Grid** tab, select the **HSM Group** tab.
2. Click the Add drop-down list and select **Thales Luna Group**.
3. In the *Add Thales Luna Group* wizard, complete the following and click **Next**:
  - **Name**: Enter a name for the HSM group.
  - **Partition Password**: Enter the partition password, and re-enter it in the **Confirm Partition Password** field.
  - **Version**: Select the Thales Luna HSM version, which is either **Luna 4**, **Luna 5**, **Luna 6**, or **Luna 7 CPL**.
  - **Comment**: You can enter additional information about the HSM.
4. Click the Add icon to add a Thales Luna HSM device, and complete the following:
  - **Name or IP Address**: Enter the hostname or IP address of the HSM device.
  - **Partition SN**: Enter the partition serial number (PSN) of the HSM. The **Partition ID** field automatically displays the ID after the configuration is saved and the appliance has successfully connected to the device.
  - **Disabled**: Select this checkbox to disable use of this HSM.
  - **Server Certificate**: Upload the certificate of the Thales Luna HSM.
5. Save the configuration.

After you add the HSM group, the Add icon and Add button in the Toolbar are greyed out. Note that if the HSM is configured in FIPS 140-2 compliant mode, certain key algorithms and key sizes are disallowed. Requests for those key algorithms or key sizes result in an error. The following algorithms are FIPS 140-2 compliant: DSA, DSA/NSEC3, RSA/SHA1, RSA/SHA1/NSEC3, RSA/SHA-256, and RSA/SHA-512. For additional information about selecting key algorithms, see [About the DNSKEY Algorithm](#).

You can verify whether the Grid Master Candidate is properly registered with the HSM by navigating to the **Grid** -> **Grid Manager** -> **Members** page. It's Status icon is yellow if it is not registered with the HSM.

If DNS service is enabled, you can also verify whether the Grid Master was able to contact the Thales Luna HSMs by navigating to the **Data Management** > **DNS** > **Members** page. If the Grid Master status is yellow, check the status of the HSMs in the **Grid** > **HSM Group** page. (For more information, see Monitoring the HSM Group below.) If the status is not green, check the configuration of the HSMs and restart the DNS service.

## Adding and Managing an Entrust nShield HSM Group

On the Entrust nShield HSM, configure the Grid Master and Grid Master Candidate as HSM clients. Enroll the IP addresses of both the Grid Master and Grid Master Candidate to avoid any service interruptions, in case the Grid Master Candidate is promoted to Grid Master. If the Grid Master and Grid Master Candidates are HA pairs, you must enroll their VIPs.



### Note

In the unlikely event that the Grid Master Candidate was registered with the Entrust nShield HSM after the Grid Master promotion, you must restart the DNS service on the newly promoted Grid Master.

In addition, you must also set up client cooperation to allow both the Grid Master and Grid Master Candidate access to the Remote File Server (RFS). Note that anytime you add a new Grid Master Candidate, you must enroll its IP address and set up a client cooperation to allow it access to the RFS. For more information on these procedures, refer to your HSM documentation.

Note that DSA cannot be used as the DNSSEC cryptographic algorithm for Entrust nShield HSMs. Therefore, migrating to Entrust nShield HSMs is not allowed if the Grid Master uses DSA as the DNSSEC cryptographic algorithm.

You can create one Entrust nShield HSM group in the Grid, and then add HSMs to the group. The appliance tries to connect to each of the HSMs in the order that they are listed.

To add an Entrust nShield HSM group:

1. From the **Grid** tab, select the **HSM Group** tab and click the Add icon.
2. In the **Add HSM Group** wizard complete the following, and then click **Next**:
  - **Name**: Enter a name for the HSM group.
  - **Protection**: Select the level of protection that the HSM group uses for the DNSSEC key data.

- **Module:** Select this if the HSM group uses a module-protected key. You do not have to enter a password phrase for this type of key.
  - **Softcard:** Select this if the HSM group uses a softcard-protected key. You must then specify the card name and password.
  - **Card Name:** Enter a name for the softcard.
  - **Password Phrase:** Enter the password and re-enter it in the **Confirm Password Phrase** field.
  - **RFS IP Address:** Enter the remote file server (RFS) IP address. Note that you must ensure that you enter a valid RFS IP address for the Security World. Validation is limited to IP address checking. Infoblox recommends that you use **Test HSM Group** to check the HSM group configuration before proceeding.
  - **RFS Port:** Specify the port of the RFS.
  - **Comment:** Optionally, enter additional information about the group.
3. To add modules to the group, click the Add icon and complete the following:
- **Remote IP:** Enter the IP address of the HSM.
  - **Remote Port:** Specify the destination port on the HSM. The firewall must be configured to allow connection to this port.
  - **Disabled:** Select this checkbox to disable use of this HSM.
  - **Keyhash:** Enter the keyhash, which is displayed on the console of the HSM. It can be obtained through an out of band mechanism from the HSM administrator. Note that the appliance validates the keyhash. If the entry is correct, the appliance displays the Electronic Serial Number (ESN) of the HSM when the editor is next launched. If the keyhash is incorrect, the appliance does not connect to the HSM.
  - **ESN:** This is a read-only field that displays the ESN of the HSM after you save the configuration and relaunch the editor. Infoblox strongly recommends that you verify the ESN displayed by the appliance with the one obtained from the HSM administrator to ensure that the appliance is communicating with the correct HSM.
4. Save the configuration.

## Monitoring the HSM Group

You can monitor the status of the HSM group and of individual modules in the group by navigating to the **Grid** tab > **HSM Group** panel. To view the status of each HSM, click the arrow beside the group name. This panel displays the following information:

- **Name:** The name of the HSM group or module.
- **Status:** The HSM group status displays the status for all the HSMs in the group. The status icon can be one of the following:
  - Green:** All the HSMs in the group are functioning properly.
  - Yellow:** At least one HSM in the group is not functioning properly.
  - Red:** All the HSMs in the group are not functioning properly.
  - Black:** The status of the HSM devices is unknown.

The status icon for each HSM can be one of the following:

- Green:** The HSM is functioning properly. For Thales Luna 5 or 6 devices, the status icon can also display **x%used** which refers to the storage capacity of the HSM partition that is assigned to the Grid. Note that when the capacity reaches 100%, new zone signings and key rollovers for existing zones cannot be performed.
- Red:** The HSM is not functioning properly. For a Thales Luna HSM, this indicates that the Grid Master was able to connect to the HSM, but no partition was assigned to the Grid Master.
- Black:** The status of the HSM device is unknown.
- **FIPS:** This applies to a Thales Luna HSM only. It indicates if the HSM is in FIPS compliant mode.
- **Comment:** Any comments that were entered about the HSM group.

You can also do the following in this tab:

- Sort the data in ascending or descending order by column.
- Print and export the data in this tab.

## Enabling HSM Signing

When you enable HSM signing, the HSM starts generating the DNSSEC keys at the next key rollover. For information about key rollovers, see [About Key Rollovers](#). You can enable this feature at the Grid level only.

To enable HSM signing:

1. From the **Data Management** tab -> **DNS** tab, expand the Toolbar and click **Grid DNS Properties**.
2. In the *Grid DNS Properties* editor, Click **Toggle Advanced Mode**, if the editor is in Basic mode, and then select the **DNSSEC** tab.
3. In the **DNSSEC** tab, select the **Enable DNSSEC** checkbox, if it is not selected, and then select the **Enable HSM Signing** checkbox.
4. Complete the other fields described in [Configuring DNSSEC Parameters](#). Note that Entrust nShield HSMs do not support DSA.
5. Save the configuration.

## Testing the HSM Group

After you configure the HSM group, you can test the HSM signing functionality of the group. Click **Test HSM Group** in the Toolbar, and then click **Yes** when the confirmation dialog displays. The appliance then executes the command to perform a signing test. The feedback panel displays the status of the test in the Grid Manager feedback panel.

## Synchronizing the HSM Group

You can click **Resync HSM Group** in the Toolbar to do any of the following:

- For an Entrust nShield HSM group, if the RFS security settings change use this function to have the appliance perform an RFS synchronization.
- For a Thales Luna HSM group, use this function to synchronize the keys of the HSM members in the group.

## Configuring Grid Members to Support DNSSEC as Secondary Servers

Any Infoblox Grid member can function as a secondary server for DNSSEC signed zones. It can receive transfers of signed zones from the Grid Master or an external primary server, and from other secondary servers. It can also respond to queries for DNS data in DNSSEC signed zones for which it is a secondary server.

### Configuring a Secondary Server for Signed Zones

The following are the tasks to configure an appliance as a secondary server for signed zones:

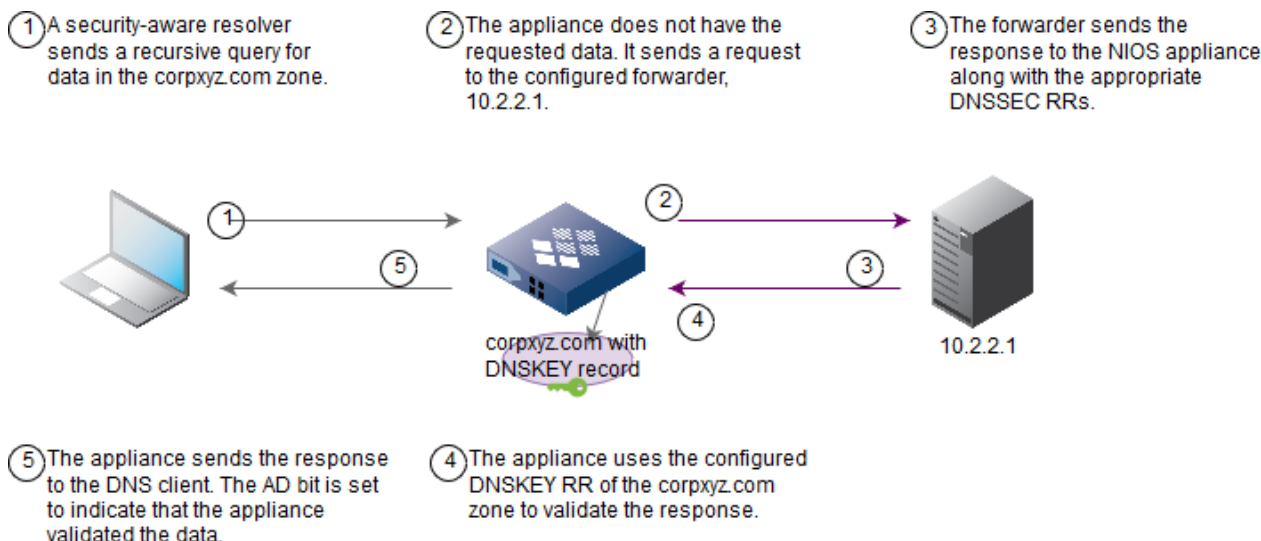
1. Enable DNSSEC on the appliance. For information, see [Enabling DNSSEC](#).
2. Configure the appliance as a secondary server for the zone. For information, see [Specifying Secondary Servers](#). If the primary server for the signed zone is external, then you must allow zone transfers to the secondary server.
3. For information, see [Enabling Zone Transfers](#). If the primary server is the Grid Master, then the secondary server receives data through the Grid replication process by default.

## Configuring Recursion and Validation for Signed Zones

When you enable recursion on a Grid member and it receives a recursive query for DNS data it does not have, it queries remote name servers that you specified in the *Grid DNS Properties* or *Member DNS Properties* editor. It then includes the DNSSEC data it retrieved through recursion in its responses to clients that requested DNSSEC RRs. You can enable the appliance to validate the responses of these servers for certain zones. On the appliance, you specify the zones to validate and configure their DNSKEY records as trust anchors. When the appliance validates a response for a zone configured with a trust anchor or for any of its child zones, the appliance starts with the DNSKEY that you configured and proceeds recursively down the DNS tree.

In the example shown in the below figure, the following was configured on the NIOS appliance:

- Forwarder with the following IP address: 10.2.2.1
- Recursion was enabled
- DNSSEC and validation were enabled
- The corpxyz.com zone and its DNSKEY record were configured



## Enabling Recursion and Validation for Signed Zones

The following are the tasks to enable recursion and validate recursively derived data:

1. Enable DNSSEC on the appliance. For information, see [Enabling DNSSEC](#).
2. Enable validation and configure the trust anchor of each signed zone. For information, see [Enabling DNSSEC Validation](#) below. You must configure at least one trusted DNSKEY RR.
3. Enable recursion on the appliance. For information, see [Enabling Recursive Queries](#).
4. Complete any of the following:
  - Configure the forward, delegated, stub or root zones for the signed zones. For information, see [Configuring Delegated, Forward, and Stub Zones](#) and [Creating a Root Zone](#).
  - Configure global forwarders and custom root name servers, if needed. For information, see [Using Forwarders](#) and [About Root Name Servers](#).

## Enabling DNSSEC Validation

To configure trust anchors and enable Infoblox name servers to validate responses:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **Members** tab -> *member* checkbox and click the Edit icon.  
**DNS View:** From the **Data Management** tab, select the **Zones** tab -> *dns\_view* checkbox and click the Edit icon.  
To override an inherited property, click **Override** next to the property to enable the configuration.
2. In the editor, click **Toggle Expert Mode**.
3. When the additional tabs appear, click **DNSSEC**.
4. In the **DNSSEC** tab, complete the following:
  - **Enable DNSSEC validation:** If you allow the appliance to respond to recursive queries, you can select this checkbox to enable the appliance to validate responses to recursive queries for domains that you specify. You must configure the DNSKEY RR of each domain that you specify.
  - **Accept expired signatures:** Click this checkbox to enable the appliance to accept responses with signatures that have expired. Though enabling this feature might be necessary to work temporarily with zones that have not had their signatures updated in a timely fashion, note that it could also increase the vulnerability of your network to replay attacks.
  - **Trust Anchors:** Configure the DNSKEY record that holds the KSK as a trust anchor for each zone for which the Grid member returns validated data. Click the Add icon and complete the following:



- **Zone:** Enter the FQDN of the domain for which the member validates responses to recursive queries.
- **Secure Entry Point (SEP):** This checkbox is enabled by default to indicate that you are configuring a KSK.
- **Responses must be secure:** Enable this checkbox to make all responses to domains configured with a trust anchor to be DNSSEC secure and valid. When you enable this checkbox, the appliance returns SERVFAIL responses for the domains that are not DNSSEC secure. Note that for each anchor, the current setting of **Responses must be secure** will be preserved when NIOS is upgraded.
- **Algorithm:** Select the algorithm of the DNSKEY record: **RSA/SHA1(5), DSA (3), DSA/NSEC3 (6), RSA/MD5 (1), RSA/SHA1/NSEC3 (7), RSA/SHA-256 (8), RSA-SHA-512 (10), ECDSAP/SHA-256 (13), or ECDSAP/SHA-384 (14)**. This must be the same algorithm that was used to generate the keys that were used to sign the zones.
- **Public Key:** Paste the key into this text box. You can use either of the following commands to retrieve the key:
  - `dig . dnskey +multiline`  
The above command retrieves root zone keys and is the only public key you require for full chain of trust validation.
  - `dig \[@server_address\] <zone> dnskey +multiline +dnssec`  
The above command retrieves public keys from the zone you specify on the server and can be used if the parent zone is not signed.  
Note that the aforementioned command provides you with a key you need to cross validate against other servers to ensure you have an identical key.  
As an alternative, you can use <http://data.iana.org/root-anchors/> to retrieve signed public keys. You can find the trust anchors in formats like XML and CSR. For more information, refer to <http://data.iana.org/root-anchors/>.
- **Negative Trust Anchors:** Configure negative trust anchors to suppress DNSSEC validation for certain domains. Click the Add icon to add the domain name to the list. You can define negative trust anchors at the Grid level and override them at the member and DNS view levels. For more information about negative trust anchors, see [Defining Negative Trust Anchors](#) below.  
To delete a negative trust anchor, select the checkbox adjacent to the **Zone** column and click the Delete icon.

5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Defining Negative Trust Anchors

A DNSSEC misconfiguration is not uncommon, and it can lead to failures in validating clients for particular domains. You can use negative trust anchors to enhance the deployment of DNSSEC validation even with misconfigured domains, as the administrator can enable DNSSEC validation without worrying about resolution failure for misconfigured signed domains that succeeds without DNSSEC validation.

A negative trust anchor is a domain name for which DNSSEC validation must be suppressed even if the domain name is listed under a trust anchor.

You can define a negative trust anchor at the Grid level and override it at the member and DNS view levels. You can define negative trust anchors within a view or at the Grid level. If you define it at both levels, only the configuration with the view will be effective for the respective view.

For any specific domain name, negative trust anchors are mutually exclusive of trusted anchors; if a negative trust anchor is specified for a domain name, you cannot configure a trust anchor using the same name. Likewise, if a trust anchor is configured for a domain name, you cannot configure negative trust anchor using the same name.

Note that if there is a trust anchor for a specific name under a negative trust anchor, that trust anchor will re-enable DNSSEC validation for query names covered by the trust anchor. For example, if you configure trust anchors for "." and "example.com," and "com" is in the negative trust anchor list, queries for "www.example.com" are subject to DNSSEC validation, while DNSSEC validation will be suppressed for www.insecure.com, even if "." is listed in the trust anchor.

The appliance does not support automatic cleanup of negative trust anchors, nor does it provide any information about their expiration. Administrators must manually keep track of the status of domains listed as a negative trust anchor, and must remove them from the list as soon as the domains become DNSSEC signed correctly. Infoblox recommends that you do not use negative trust anchors, and rather disable DNSSEC validation, if the administrator is not familiar with negative trust anchors or is not able to maintain the negative trust anchors properly; careless use of negative trust



anchors would rather hinder the deployment of DNSSEC, which is the opposite to the purpose of negative trust anchors. For more information about general technical details of negative trust anchors, refer to <http://tools.ietf.org/html/draft-livingood-negative-trust-anchors>.

Note the following about negative trust anchors:

- You must restart the DNS service when you modify the list of negative trust anchors.
- The appliance displays an error message if an entry is present in both the trust anchors and the negative trust anchors list for the same FQDN.
- The appliance displays an error message if the same FQDN is present multiple times in a negative trust anchor.
- When DNSSEC validation is suppressed due to a negative trust anchor, the corresponding response from the validating resolver does not include the AD bit.

## Applying Policies and Rules to DNS Queries that Request DNSSEC Data

You can configure the NIOS appliance to always apply RPZ policies, DNS blacklists, or NXDOMAIN rules to DNS queries, regardless of whether the queries request DNSSEC data. You can also configure the appliance to generate synthesized AAAA records for DNS queries that request DNSSEC data.



### Warning

*When you enable this feature, NIOS applies the selected policies and rules even when it responds to DNS clients that support DNSSEC.*

*Note that responses to these clients may result in resolution failure. Infoblox recommends that you use caution when enabling this feature and DNSSEC validation at the same time.*

To enable this feature, complete the following:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab. Expand the Toolbar and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox and click the Edit icon.  
**DNS View:** From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *default dns\_view* and click the Edit icon.
2. In the editor, click **Toggle Expert Mode**, if the editor is in Basic mode, and then select the **DNSSEC** tab.
3. In the **Apply the selected policies/rules to queries requesting DNSSEC records** section, complete the following:
  - **Response Policy Zones (RPZ) policies:** Select this to apply RPZ policies to DNS queries that request DNSSEC data. You must install the RPZ license to view this checkbox. For more information, see [About Infoblox DNS Firewall](#).
  - **Blacklist rules:** Select this to apply blacklist rules to DNS queries that request DNSSEC data. For more information, see [About Blacklists](#).
  - **NXDOMAIN rules:** Select this to apply NXDOMAIN rules to DNS queries that request DNSSEC data. This checkbox is visible only if you install the Query Redirection license. For more information, see [About NXDOMAIN Redirection](#).
  - **DNS64 Groups:** Select this to generate synthesized AAAA records for DNS queries that request DNSSEC data. For more information, see [Enabling DNS64 Service](#).  
The member and the DNS views inherit settings from the Grid, by default. To override the settings, click **Override**. You can override settings at the member and DNS view levels. To retain the same settings as the Grid, click **Inherit** at the member and DNS view level.
4. Save the configuration.

## Managing DNS Traffic Control

This section describes the Infoblox DNS Traffic Control solution. It provides guidelines about how to adjust DNS responses based on DNS query source IP, server availability, and network topology. It also describes the required objects in this process such as DTC (DNS Traffic Control) servers, pools, and LBDNs (Load Balanced Domain Names), and how to configure them to achieve the best load balancing results.

It includes the following topics:

- [About DNS Traffic Control](#)
- [License Requirements and Admin Permissions](#)
- [Limitations of DNS Traffic Control](#)
- [Load Balancing Methods for DNS Traffic Control](#)
- [Configuring Topology Rules and Rulesets](#)
- [Managing DNS Traffic Control Objects](#)
- [Visualization for DNS Traffic Control Objects](#)
- [Backing Up DTC Configuration Files](#)
- [Using DNS Traffic Control Health Monitors](#)
- [Configuring DTC Monitors for Health Check](#)

## About DNS Traffic Control

Infoblox DNS Traffic Control (DTC) provides a load balancing solution by creating DNS responses based on DNS query source IP, server availability, and network topology. Through DTC, you can set up multiple global sites and configure supported objects and load balancing methods to direct DNS clients to the best available servers. For detailed information about how DNS Traffic Control handles DNS queries and responses, see [DNS Traffic Control Querying Process](#) below.

To use the DTC feature and bypass the standard DNS querying process, you must install the DNS Traffic Control license on designated Grid members. For information about license and admin requirements, see [License Requirements and Admin Permissions](#). Members that are not authoritative for zones or members that do not have the DTC license installed will not process DNS queries through DTC. However, the appliance can process DNS queries through DNS Traffic Control for secondary servers in the Grid using the data replication method.

DNS Traffic Control utilizes a load balancing mechanism to create DNS responses. It returns tailored DNS responses based on settings you configure for associated objects such as DTC servers, pools, and LBDNs. For more information about these objects, see [Supported DNS Traffic Control Objects](#) below and [Managing DNS Traffic Control Objects](#). You can configure load balancing methods for pools and LBDNs based on the source IP address and other criteria. For more information, see [Load Balancing Methods for DNS Traffic Control](#).

To ensure that DTC servers are reachable and can process DNS queries, you can configure health monitors that help you determine the availability of these servers. For more information, see [Using DNS Traffic Control Health Monitors](#).

After you have set up DNS Traffic Control for specific DTC objects, you can monitor their status as described in the section [Viewing DNS Traffic Control Objects](#). You can also view a visualization of the hierarchy of DNS Traffic Control objects that you configured. For more information, see [Visualization for DNS Traffic Control Objects](#).

You can enable or disable logging for DNS Traffic Control load balancing and health monitors. The appliance logs this information to the syslog. For more information, see [Setting DNS Logging Categories](#).

You can configure the DNS Traffic Control properties for the Grid and Grid members. For more information, see [Configuring DNS Traffic Control Properties](#).

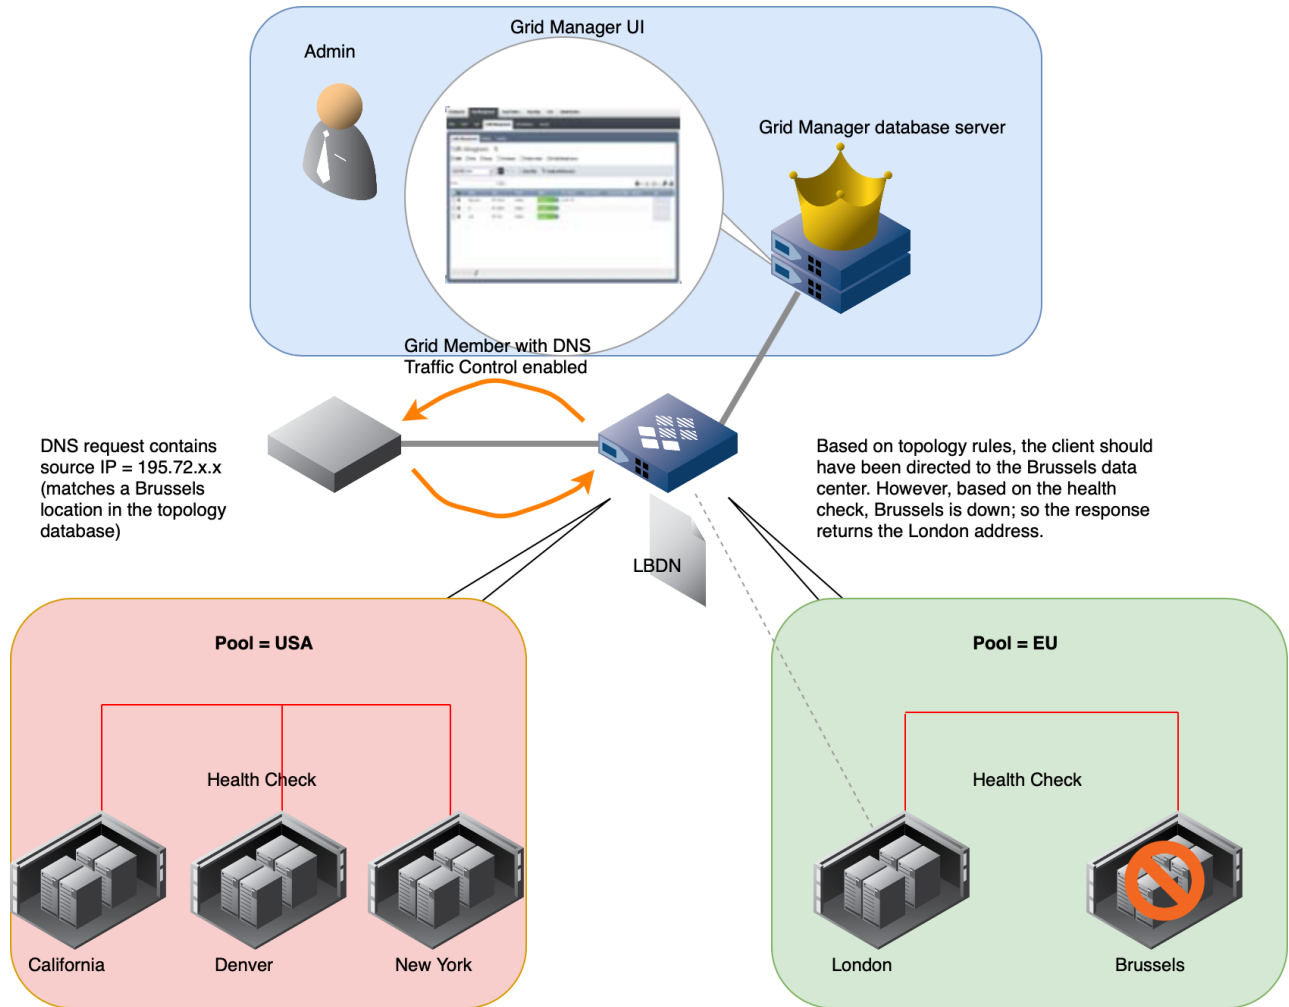
## DNS Traffic Control Configuration Example

The *DNS Traffic Control Example* figure below illustrates the basic concept of DNS Traffic Control and how you can use it to direct DNS clients to the best available server using the Topology load balancing method. In the illustration, consider your company manages five data centers: California, Denver, New York, London, and Brussels. Each data center offers the same services and applications your users need. To optimize server utilization and minimize response time, you use DNS Traffic Control to provide DNS responses based on the source IP address and the geographic locations of your data centers. You define your data centers as DTC server objects, and then add them to a pool based on their locations. Data centers in California, Denver, and New York go into Pool USA while London and Brussels go into Pool EU. You also create a topology ruleset containing geography rules that is used when you configure Topology as the load balancing method for your LBDN and pools. To ensure that your data centers are available, you also configure health monitors so that your pools can check the health of its DTC servers.

In DNS Traffic Control, you complete the following to set up the above configuration:

1. Define a geography rule for the topology ruleset, as described in [Configuring Topology Rules and Rulesets](#).
2. Configure each data center as a DTC server object, as described in [Configuring DNS Traffic Control Servers](#).
3. Configure health monitors that NIOS will use to check the availability of DTC servers, as described in [Using DNS Traffic Control Health Monitors](#).
4. Create two pools (USA and EU), which group your servers by geographical locations, configure health monitors that you created earlier, and then select Topology as the load balancing method. For information, see [Configuring DNS Traffic Control Pools](#).
5. Define an LBDN and select Topology as the load balancing method and then associate it with a DNS zone. The appliance can then match the geography of the source IP addresses and direct the client to the best available server. For information, see [Configuring DNS Traffic Control LBDNs](#).

*DNS Traffic Control Example*



To achieve load balancing results for DNS Traffic Control, you can configure DTC objects in the following order:

1. Create DNS Traffic Control servers for each data center or server you want to manage. For information, see [Configuring DNS Traffic Control Servers](#).
2. Optionally, if you want to monitor server health, configure health monitors and add them to your pools when you create them. For information, see [Using DNS Traffic Control Health Monitors](#).
3. Configure any topology rulesets that will be used by DTC pools. For information, see [Defining Topology Rulesets](#).
4. Configure DTC pools, as described in [Configuring DNS Traffic Control Pools](#).

5. Configure any topology rulesets that will be used with LBDNs, as described in [Configuring Topology Rules and Rulesets](#).
6. Configure DTC LBDNs, as described in [Configuring DNS Traffic Control LBDNs](#).

## DNS Traffic Control Querying Process

DNS Traffic Control handles a DNS query if the query name belongs to a zone for which the appliance is authoritative and matches an LBDN that is linked to the respective zone. Otherwise, the appliance processes DNS queries normally using the standard DNS query processing.

If you have configured persistence for DTC LBDN and the DTC cache contains a previous response for the same client, then DTC returns the cached response to the respective client. Otherwise, the DNS Traffic Control querying process first selects an LBDN, a pool, and then a specific server from that pool. A DNS record is synthesized from the address of the selected server and returns only DTC LBDN records to the client. Note that the configured DNS records are obscured by the DTC LBDN records. The obscured records are indicated by a strikethrough. For example, an obscured A record appears as A Record in Grid Manager.

When all the pools associated with the LBDNs are unavailable, the appliance returns DNS data for the zone. For information about how to configure availability and topology rules, see [Configuring Topology Rules and Rulesets](#). Following is a brief description of the DNS Traffic Control querying process:

1. The DNS Traffic Control querying process first checks an LBDN's DTC cache to verify if a response is available for the same client, same query and if the server in the cached response is online. If these conditions match, it returns the cached response to the client and refreshes the cache expiry time.
2. If the response is not available in the DNS Traffic Control cache, the DNS Traffic Control querying process occurs as follows:
  - Based on the FQDN in the request, the DNS Traffic Control querying process selects a matching LBDN based on its associated zone and pattern.
  - Based on the selected LBDN load balancing method, the DNS Traffic Control querying process selects an available pool. It does not select pools that are not available or do not have online servers associated with it. If pools are not available for the selected LBDN, the DNS Traffic Control querying process fails to determine the result.
  - After selecting a pool, the DNS Traffic Control querying process selects a server from the pool, based on the load balancing method that you have selected for the respective pool. The DNS Traffic Control querying process uses the preferred load balancing method of the pool by default and the alternate method only if the preferred method fails to return a result.
  - If the DNS Traffic Control querying process fails to determine a result, then the DNS server responds to the client with the matching DNS records configured for the respective zone. If matching DNS records are not configured, it returns an empty result. You can enable or disable the DTC to DNS fall through in the Grid DNS properties.

## Supported DNS Traffic Control Objects

You can configure the following DNS Traffic Control objects in the Grid:

- Server: DNS Traffic Control synthesizes DNS records for the servers. For information about how to add and modify servers, see [Configuring DNS Traffic Control Servers](#).
- Pool: A pool is a collection of servers that you can put together as possible responses to queries. For information about how to add and modify load balancing pools, see [Configuring DNS Traffic Control Pools](#).
- LBDN: A DTC LBDN is a load balanced domain name object that is used by DNS Traffic Control to process DNS queries for load balanced resources. For more information about how to add and modify LBDN objects, see [Configuring DNS Traffic Control LBDNs](#).

## License Requirements and Admin Permissions

The DNS Traffic Control works as an add-on feature for the DNS service in the NIOS appliance. In order to use it, install the DNS and DNS Traffic Control licenses on the appliance. You can add the DNS Traffic Control license only when you have installed the DNS license. After you add the license, the feature becomes available to you in **Data Management->DNS -> Traffic Control**.

The DNS Traffic Control feature starts, stops, or restarts with the DNS service. You may need to restart the DNS service

after you make configuration changes. When you click **Restart** at the top of the screen, the DNS Traffic Control service takes some time to update the configuration and the status of LBDNs, pools and servers on the appliance. For more information about how to install licenses, see [Managing Licenses](#).

The appliance creates a new role, **DTC Admin**, when you install the DNS Traffic Control license. For more information about how to define permissions, see [About Admin Roles](#).

## License Requirements for DNS Traffic Control

- If you are using HA for DNS, Infoblox best practice recommendation is that you apply the DNS Traffic Control license on both the active and passive nodes of the HA pair. If the DNS Traffic Control license is applied only on the active node but not the passive node, the DNS query that matches with LBDN will elicit a DNS Traffic Control response. However, if the DNS Traffic Control license is applied only on the passive node and not the active node, then the DNS query will not elicit a DNS Traffic Control response. You must pass the DNS query using only the virtual IP address.
- A DNS Traffic Control license is required for the name servers in a load balanced zone. If there is no DNS Traffic Control license, then that server will not take part in DNS Traffic Control.

## Admin Permissions for DNS Traffic Control

- You can update, add, or delete LBDNs if you have read/write permission on **All DTC LBDNs**, but you need additional permissions to link them to a zone.
- You must have read/write permission on **All DTC LBDNs** to modify an LBDN if a configuration change results in a deletion of LBDN records.
- You must have read/write permission on both **All DTC LBDNs** and **All DTC LBDN Records** to modify LBDN patterns that might result in the creation of new LBDN records.
- Administrators who have access to set up zones have read/write access to their LBDN records. The LBDN records permission is similar to other zone record permissions.
- When you link an LBDN to a zone, you must have both read-only permission on **All DTC LBDNs** and read/write permission on the LBDN record of the zone.
- You must have read/write permission on **All DTC LBDNs** to disable an LBDN. This implicitly disables associated LBDN records from the associated zones and does not require an LBDN record permission.
- You must have read/write permission on **All Topology Databases** to upload a MaxMind GeoIP location database. To modify a rule in the topology ruleset, you must have read/write permission on **All DTC Topologies**.
- To create and modify DTC pools, you must have read/write permission on **All DTC Pools**.
- You must have read/write permission on **All DTC Servers** to create and modify DTC servers.
- You must have read/write permission on **All DTC Monitors**, **All DTC Certificates**, and **All DTC Topologies** to create and modify DTC monitors, certificates and topologies respectively.

For more information about defining global permissions, see [Defining Global Permissions](#).

## Limitations of DNS Traffic Control

- A member will return DNS Traffic Control results for a zone only if the member is a Grid primary, or a Grid secondary that is using database replication. DNS Traffic Control results are not produced for zones using AXFR, regardless of whether or not the primary member is in the Grid.
- The DNS Traffic Control querying process is not supported for recursive queries.
- The DNS Traffic Control does not support the Global application of an LBDN pattern against all queries. The appliance returns a result only if the query resolves to an authoritative zone to which an DNS Traffic Control LBDN is explicitly linked.
- You can use the **IDN converter** for conversion, but the appliance supports manually encoded puny code as an LBDN pattern.
- DNS Traffic Control health monitoring does not monitor servers if the server has IPv4 and/or IPv6 addresses and the health monitoring interface on the Grid member does not have the corresponding IP address type. The appliance sets the status as unknown for such IPv4 and IPv6 addresses. In addition, the appliance may return a timeout error while loading the Traffic Control tab in Grid Manager if you have configured health monitoring for a lot of DTC servers.

## Load Balancing Methods for DNS Traffic Control

You can define the following load balancing methods:

Load Balancing Method	LBDN	Pool
All Available	-	+
Global Availability	+	+
Source IP Hash	+	+
Round Robin	+	+
Ratio: Fixed	+	+
Ratio: Dynamic (Round Trip Delay)	-	+
Ratio: Dynamic (SNMP)	-	+
Topology	+	+

Based on the load balancing method defined for an LBDN, the DNS Traffic Control selects an available pool. Based on the method selected for a pool, it selects an available server.

The following is a description of the load balancing methods with examples for pools or LBDNs:

- Using the **All Available** method, the appliance responds to the query with all the available servers in the DTC pool for the appropriate record type. The responses are returned in the same order in which the servers are listed in the DTC pool, eliminating the unavailable servers.

Consider the following example for all available records load balancing method with an LBDN x.abc.com:

Pool= Pool\_1; load balancing method= all available records; health check= HTTPS

10.10.1.1; Availability = up

2001:db8:a22:a00::29; Availability = up

10.10.2.2; Availability = down

2001:db8:a22:a00::32; Availability = up

10.10.3.3; Availability = up

In this example, the appliance responds with 10.10.1.1, 10.10.3.3 for each A record query. For each AAAA record query, NIOS responds with 2001:db8:a22:a00::29, 2001:db8:a22:a00::32. The unavailable servers are eliminated. Note that the system considers only the order of the servers in the DTC pool and ignores the weight of available servers.

- Using the **Global Availability** method, the appliance creates the response to the query, so that the clients are directed to the first available pool and server.

Consider the following example for global availability load balancing method with an LBDN x.abc.com:

- Pool= Pool\_1; load balancing method= global availability; health check= HTTPS 10.10.1.1

10.10.2.2

10.10.3.3

- Pool= Pool\_2; load balancing method = global availability; health check = HTTPS 10.10.4.4

10.10.5.5

10.10.5.5

In this example, the appliance always responds with 10.10.1.1 (A record) for all x.abc.com queries assuming that all servers are available. The DNS Traffic Control LBDN determines which pool to use based on the health check

response and adjusts the response accordingly. The appliance responds with 10.10.4.4 from Pool\_2 if health check for all servers associated with Pool\_1 fails.

- Using the **Source IP Hash** method, NIOS matches an IP address from an incoming query with the health statuses of pools and servers to address the responses by the client. When multiple pools or servers are configured, NIOS uses the source IP hash load balancer method, a load-balancing pattern in which requests are distributed based on the hash value of an IP address from an incoming query and the health status of the pool or server. With the source IP hash load balancing method, clients have their own pool or server and are always associated with the same pool or server for the same query as long as the pool or server is green; if the health status of the pool or server turns red, NIOS switches the clients to the working pool or server and switches back when the health restores to green.

Consider the following example for the source IP hash load balancing method. The DNS Traffic Control pool has two servers and uses the source IP hash load balancer. When the health statuses are green for both the servers, depending on the client IP address and DNS Traffic Control health status, the pool returns the DNS record from the first DNS Traffic Control server for the first client and returns the DNS record from the second DNS Traffic Control Server for the second client. This continues as long as the health statuses are green; if the health status turns red for a DNS Traffic Control server, the pool switches the client to the working server and switches it back when the health status restores to green.

For information on the Limitations and Warnings of Source IP Hash Load Balancing Method, see the following section, Limitations of Source IP Hash Load Balancing Method.

- Using the **Ratio: Fixed** method, NIOS adjusts the response to the query so that the clients are directed to servers in a pool or among pools. When multiple pools or servers are configured, the appliance uses the weighted round robin method, a load-balancing pattern in which requests are distributed among several pools or servers based on a weight assigned to each pool or server. Note that the system considers the weight of available servers only. Consider the following example for ratio load balancing method with an LBDN x.abc.com, load balancing method = **Ratio** and two linked pools: Pool\_1 with weight = 70 and Pool\_2 with weight = 30.
  - Pool = Pool\_1; load balancing method = ratio; health check = HTTPS 10.10.1.1; Weight = 50  
10.10.2.2; Weight = 2  
10.10.3.3; Weight = 25
  - Pool = Pool\_2; load balancing method = ratio; health check = HTTPS 10.10.4.4; Weight = 50  
10.10.5.5; Weight = 25  
10.10.5.5; Weight = 25

In this example, the appliance responds 70% of the time with a server associated with Pool\_1. Within Pool\_1, it responds with 10.10.1.1 address 50% of the time.

- Using the **Ratio: Dynamic** method, the appliance weights the DTC servers dynamically based on round trip delay or SNMP health monitor data. You can use one of the following options:
  - **Round trip delay:** Based on the round trip delay from the DTC member that received a client's DNS request, the system sends clients to the server with the minimal latency time, i.e. the closest one. You need a pre-configured health monitor for this load balancing method.

For example:

- Server A latency = 25 ms
- Server C latency = 18 ms
- Server D latency = 50 ms
- In this case, the traffic distribution is as follows:
  - Server A = 0%
  - Server C = 100%
  - Server D = 0%

- **SNMP:** Based on data from the SNMP monitor associated to the server, for example, CPU or memory utilization, the system sends clients to the server with the lowest load. For this load balancing method, you need a pre-configured SNMP health monitor with a required metric to be tracked. The metric is set through an object identifier (OID) in the monitor properties. This method supports only OIDs for which the server can return an integer value.

The value of the monitored metric defines how the traffic is directed. By default, the servers with the highest metric values receive the client requests. There may be cases when your selected metric reflects server availability in the opposite way, that is, the lowest metric values indicate available servers. For such cases, you can invert the value of the OID, that is, of the monitored metric, and have the traffic directed to the lowest-rated servers.

You can select to weight servers by either priority or ratio. In case of priority, traffic is directed towards the servers



that report the best metric values, other servers being bypassed. In case of ratio, traffic is distributed across all servers based on the values of the monitored metric for each of them. If a health check for a server is failed, the server is excluded from the load balancing.

Consider the following example where the CPU utilization metric is used for server monitoring:

- Server A CPU utilization = 90%
- Server C CPU utilization = 50%
- Server D CPU utilization = 10%

With normal dynamic weights, the distribution is as follows:

- Server A 60% (calculated as  $90 / (10 + 50 + 90) = 90/150 = 0.6$ )
- Server C 33% (calculated as  $50 / (10 + 50 + 90) = 50/150 \sim 0.33$ )
- Server D 7% (calculated as  $10 / (10 + 50 + 90) = 10/150 \sim 0.066$ )

This means that the most loaded server will receive most requests than the less loaded one. For this case, the metric should be inverted to reflect server availability appropriately:

- Server A 8% (calculated as  $1/90 / (1/90 + 1/50 + 1/10) = 0.011/0.131 \sim 0.08$ )
- Server C 15% (calculated as  $1/50 / (1/90 + 1/50 + 1/10) = 0.02/0.131 \sim 0.15$ )
- Server D 77% (calculated as  $1/10 / (1/90 + 1/50 + 1/10) = 0.1/0.131 \sim 0.77$ )



#### Note

You can see traffic distribution percentage across members in pools and servers based on selected load balancing methods in the visualization panel. For information, see [Visualization for DNS Traffic Control Objects](#).

- Using the **Round Robin** method, the appliance returns servers sequentially and cyclically. Consider the following example for round robin load balancing method with an LBDN x.abc.com:  
Pool = Pool\_1; load balancing method = Round Robin; health check = HTTPS 10.10.1.1;  
10.10.2.2;  
10.10.3.3;

In this example, NIOS responds with a server associated with Pool\_1. Within Pool\_1, it responds sequentially:

Time 1: Response = 10.10.1.1  
Time 2: Response = 10.10.2.2  
Time 3: Response = 10.10.3.3  
Time 4: Response = 10.10.1.1  
Time 5: Response = 10.10.2.2  
Time 6: Response = 10.10.3.3

- Using the **Topology** method, the appliance applies a predefined geographic mapping method and evaluates user-defined geography, subnet, or extensible attribute rules sequentially. Geographical locations for the geography rules are provided through an external topology database. The appliance supports the MaxMind GeoIP2 City or Country database and the MaxMind GeoLite2 City or Country database. For more information, see the following section, [Configuring Topology Rules and Rulesets](#).

#### Limitations of Source IP Hash Load Balancing Method

- The source IP hash load balancing method does not balance clients evenly between DNS Traffic Control pools and DNS Traffic Control servers because the load is divided based on the IP addresses of the clients and the health statuses of the DNS Traffic Control servers.
- For the source IP hash load balancing method with the **Auto Consolidated Monitor** option enabled, when the health status of a DNS Traffic Control server changes and simultaneously two DNS Traffic Control Grid members receive the same DNS requests from a single client, the response from the Grid members may vary as the DNS Traffic Control consolidated monitors need a little time to share health results across DNS Traffic Control Grid members.





#### Warning

- If you use the source IP hash load balancing method without enabling the **Auto Consolidated Monitors** option, a warning message to enable the option is displayed to synchronize health statuses between all DNS Traffic Control Grid members. This synchronization helps in getting the same DNS replies for each DNS Traffic Control Grid member for one DNS request.
- NIOS cannot guarantee persistence in DNS responses for each DNS Traffic Control Grid member when a part of the DNS Traffic Control configuration with the source IP hash load balancing method has an active DNS Traffic Control DNS cache (LBDN persistence more than 0) even after enabling the **Auto Consolidated Monitors** option. DNS Traffic Control Grid members may respond with different DNS Traffic Control records for one DNS request because DNS Traffic Control caches can have differences between DNS Traffic Control Grid members for the same DNS request.
- When a part of the DNS Traffic Control configuration with the **Auto Consolidated Monitors** option enabled and the source IP hash load balancing method has a DNS Traffic Control server with health monitors assigned to it, a warning message asking you to manually switch the health monitors from the DNS Traffic Control server to a DNS Traffic Control pool is displayed.

## Configuring Topology Rules and Rulesets

A topology rule maps a client IP address to a DNS Traffic Control (DTC) pool or server. To use Topology as the load balancing method for a pool or an LBDN, you must define a topology ruleset containing at least one rule. The rulesets are configured globally. When the DNS Traffic Control returns a response, it evaluates the list of rules in the topology ruleset in order and uses the first match with an available destination. The method fails if there are no matches.

You can define the following topology rules in a ruleset:

- Extensible Attribute rule
- Subnet rule
- Geography rule

In the DTC Topology ruleset for Subnet rule, Geographical rule, and Extensible Attribute topology (EA) rules there are options to choose the **NOERR/NODATA** response or the **NXDOMAIN** response. It will also allow you to set the **Destination** as **SERVER** or **POOL** for the subnet from the IPAM object. The destination for a topology ruleset is either a server or a pool. An LBDN can use only topology rulesets with a pool as the destination. A pool can use only topology rulesets with a server as the destination. You can also use CSV import to import rules into NIOS.

To use DTC topology with extensible attribute rules, you must select the extensible attributes that support DTC balancing in the *Grid DNS Properties* editor. You can select up to four extensible attributes and the order in which you select them is required for filling the extensible attributes topology rule. For example, if a ruleset is to have three extensible attributes in the order A, B, and C, while creating the extensible attributes topology rule, you must select the value of A first, then B, and finally the value of C in the *Grid DNS Properties* editor. For more information, see *Configuring Grid DNS Traffic Control Properties*.

The extensible attributes topology rules work based on the IPAM configuration. The rules are supported only on IPv4 and IPv6 networks, and are not supported on network containers and hosts. For a correct DTC load balancing based on IPAM networks, you must assign extensible attributes that are selected for DTC and have valid values for IPv4/IPv6 Networks. You can also configure extensible attributes to be inherited from parent IPv4/IPv6 network containers. For more information, see [Managing Extensible Attributes](#).

DTC topology rulesets with extensible attribute rules use a separate database instance — DTC EA database that should be up to date to ensure correct navigation for the incoming DTC queries. Each time you make changes to extensible attributes in the IPAM object, or manage a set of extensible attributes for DTC in Grid Manager, a notification prompting you to rebuild the DTC EA database is displayed. If you choose to ignore the prompt, then all changes to DTC balancing behavior that require a database rebuild will not take effect. For more information about rebuilding the EA database, see [Rebuilding EA Database](#) below.



#### Note

- If the DNS service is already running, you must not restart the service after the rebuild of the EA database. Instead, wait until the DTC members finish making the new EA database build active.
- Once the DNS Traffic Control restores, you must manually rebuild the DNS Traffic Control EA database.

#### Limitations of Configuring Topology Rules and Rulesets

- Based on the destination type of the DTC Topology ruleset settings, you can set the **NOERR/NODATA** or the **NXDOMAIN** response for a new DTC Topology rule. However, you cannot choose the destination **DTC Pool** or the **DTC Server**.
- If the **Destination Type** is set as **SERVER**, then the topology rule set for DTC pool cannot have only **NOERR** and **NXDOMAIN** rules. This is because the NAMED does not process the queries in the IDNS if the incoming requests are matched to LBDN with no topology balance method. Also, the pools under this LBDN do not have active or existing servers. Hence a **NOERR** response is always received. To prevent this behavior, the grid must have at least one active server as the **Rule Destination** under any DTC pool in the LBDN to allow IDS processing for the current LBDN and Pool. The rule set must have at least one rule with the **REGULAR** Return Type.
- The **Disable for DHCP** checkbox must not be selected to enable networks and to build the topology database. For more information about the checkbox, see [Configuring IPv4 Networks](#).
- The topology ruleset must have a specific order for the following rules:
  - a. **REGULAR** rules
  - b. **NOERR** rules
  - c. **NXDOMAIN** rules

The ruleset cannot have rules with a **REGULAR** return type after the **NOERR** or **NXDOMAIN** rules are set in order. You will receive a warning message when you try to save the topology ruleset in any other order. However, the GRID automatically sorts the rules in the correct order, once you accept the warning message.



#### Note

During the WAPI call, if the rules are not in the correct order, they are automatically sorted as WAPI does not give any warnings.

#### Defining Topology Rulesets

A topology ruleset can contain multiple rules. The rules in a topology ruleset must use the same destination type. Multiple LBDNs or pools can reuse a topology ruleset.

Each server that you use as a destination in the topology must exist in every pool that is using the topology. When you select Topology as a load balancing method for a pool, you can select one of these rulesets for the topology rules. The ruleset can be a combination of extensible attribute, subnet, and/or geography rules.

Note the following about extensible attribute, subnet, and geography source matches:

- A rule with an extensible attribute source matches if a client query comes from the network that has the specified set of extensible attributes. In other words, extensible attributes you specify when you create a rule.
- A rule with a subnet source matches if the subnet contains the client IP address.
- A rule with a geography source label matches if the client IP address and geography source label match corresponding information in the MaxMind location database.

Note the following information about rules and rulesets:

- When you upload a new MaxMind location database or restore a backup, the appliance does not automatically remove rules that contain invalid labels. Instead, it marks the rules with labels that do not exist in the database as invalid. The appliance ignores these rules during the querying process, and you cannot save the configuration if it is modified, but you can use the existing configuration.

- The appliance checks specific combinations of labels when the rules use multiple conditions. For example, if you have a rule with the source types Country = Canada and City = Vancouver and you change the Country source type to Russia, the City source type is cleared and the selector resets to contain only known cities in Russia. This is applicable for both geography and extensible attribute rules.

The following is an example of valid source types:

Continent	Country	Subdivision	City
Any	Canada	Any	Vancouver
Any	Any	Any	Vancouver
North America	Any	Any	Vancouver
North America	USA	Washington	Vancouver

- When rules have multiple source conditions, the client must match all conditions for the rule to execute.
- A ruleset may have multiple subnet rules and the subnets may overlap. Similarly, a ruleset may have multiple geography rules and the matches may overlap. Similarly, a ruleset may have multiple extensible attribute rules and the matches may overlap. During the querying process, the rules in a topology ruleset are evaluated in order. For example, if you configure subnet rules where #1 is 10.10.0.0/16 and #2 is 10.0.0.0/8, both are considered valid in the appliance.

To define a ruleset, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, and then click **Manage Topology Rulesets** in the Toolbar.
2. In the *Topology Manager* window, click the Add icon.
3. In the *Ruleset* wizard that appears, complete the following:
  - **Name:** Enter a name for the ruleset.
  - **Destination Type:** Select a destination type, **Pool**, or **Server**. Rulesets with the Pool destination type can only be used by LBDNs. Rulesets with the Server destination type can only be used by pools. You cannot change the destination type if the ruleset contains any rules.
  - **Comment:** Enter additional information about the ruleset.
  - **Rules:** You can define multiple extensible attribute rules, subnet rules, and geography rules in the ruleset. Click the arrow next to the Add icon and select either **Extensible Attribute Rule**, **Subnet Rule**, or **Geography Rule**.
    - When you select **Extensible Attribute Rule**, the Grid Manager displays the following:
      - **Source Type:** Define up to four extensible attributes to use as the source type for the EA topology ruleset. The values for the IPAM object EAs are provided from the DNS Traffic Control EAs selected in the Grid DNS Properties editor. To define extensible attribute source types for the topology rules, see *Configuring Grid DNS Traffic Control Properties*.  
Note that "Any" matches any value. There must be at least one source type with a specific value (the value is not "Any").  
When a source type uses "does not equal" as the operator, it must be the lowest level source type (most specific). For example, with Continent/Country/Subdivision/City, City is the most specific source type.
      - **Destination/Response:**
        - **DTC Pool/Server:** Click **Select** to select a destination. The appliance displays the *DTC Pool Selector* dialog box when you have selected the Pool destination type, and displays *DTC Server Selector* dialog box when you have selected the Server destination type. Click a specific pool or server to select it. Note that if there is only one pool or server, no dialog box is displayed when selecting the destination.
        - **NOERROR/NODATA (Response):** Select this option to provide a **NOERROR/NODATA** response for DTC queries.

- **NXDOMAIN (Response):** Select this option to provide an **NXDOMAIN** response for DTC queries.  
Click **Add** to add the source. The appliance displays the following information in the Rules table:
    - **Source:** The values of extensible attributes that you specified.
    - **Destination:** The destination that you selected.
    - **Valid Source:** After you save the ruleset, the value is set to **Yes** if the extensible attributes exist in the EA database.  
Note that the source must be valid when creating a ruleset. It can become invalid when a new topology database no longer contains the source.
    - **Order:** Displays the order of the rule in the ruleset.
    - **Return Type:** The response type that is selected.
- When you select **Subnet Rule**, the Grid Manager displays the following:
  - **Source Subnet:** Select a value from the drop-down list. You can either select **equals** or **does not equal**, and specify a subnet IP address or click **Select** and choose a network from the *Network Selector* dialog box.  
Note that "Any" matches any value. There must be at least one source subnet with a specific value (the value is not "Any").  
When a source subnet uses "does not equal" as the operator, it must be the lowest level source subnet (most specific).
  - **Destination/Response:**
    - **DTC Pool/Server:** Click **Select** to select a destination. The appliance displays the *DTC Pool Selector* dialog box when you have selected the Pool destination type and displays the *DTC Server Selector* dialog box when you have selected the Server destination type. Click a specific pool or server to select it. Note that if there is only one pool or server created, no dialog box is displayed when selecting the destination.
    - **NOERROR/NODATA (Response):** Select this option to provide a **NOERROR/NODATA** response for DTC queries.
    - **NXDOMAIN (Response):** Select this option to provide an **NXDOMAIN** response for DTC queries.  
Click **Add** to add the source. The appliance displays the following information in the Rules table:
      - **Source:** The subnet address that you specified.
      - **Destination:** The destination that you selected.
      - **Valid Source:** For a subnet rule, the rule is always marked as valid after you save the ruleset.
      - **Order:** Displays the order of the rule in the ruleset.
      - **Return Type:** The response type that is selected.
- When you select **Geography Rule**, Grid Manager displays the following:
  - **Source Type:** Select a source type.
  - **Continent:** Select a continent from the drop-down list. You can also enter the first few characters of the continent to match an item in the database.
  - **Country:** Select a country from the drop-down list. You can also enter the first few characters of the country to match an item in the database.
  - **Subdivision:** Select a subdivision from the drop-down list. You can also enter the first few characters of the subdivision to match an item in the database.
  - **City:** Select a city from the drop-down list. You can also enter the first few characters of the city to match an item in the database. The drop-down list has paging controls to page through the available values.
  - **Destination/Response:**
    - **DTC Pool/Server:** Click **Select** to select a destination. The appliance displays the *DTC Pool Selector* dialog box when you have selected the Pool destination type and the displays *DTC Server Selector* dialog box when you have selected the Server destination type. Click a specific pool or server to select it. Note that if there is only one pool or server created, no dialog box is displayed when selecting the destination.
    - **NOERROR/NODATA (Response):** Select this option to provide a **NOERROR/NODATA** response for DTC queries.

- **NXDOMAIN (Response)**: Select this option to provide an **NXDOMAIN** response for DTC queries.

Click **Add** to add the source. The appliance displays the following information in the Rules table:

- **Source**: The subnet address that you specified.
  - **Destination**: The destination that you selected.
  - **Valid Source**: After you save the ruleset, the value is set to **Yes** if the labels exist in the MaxMind location database.
  - **Order**: Displays the order of the rule in the ruleset.
  - **Return Type**: The response type that is selected.
  - **Default destination if none of the above rules match (optional)**: Click **Select** to select the default destination if none of the above rules match. The appliance displays the *DTC Pool Selector* dialog box when you have selected the Pool destination type and displays the *DTC Server Selector* dialog box when you have selected the Server destination type. Click a specific pool or server to select it. You can click **Clear** to remove the selected pool or server. Note that you can select a default destination even if there are no rules defined in the Rules table.
4. If necessary, modify the order of rules in the table. You can do so by editing the value in the **Order** column or by using the arrows on the left-hand side of the table.
  5. Click **Next**.
  6. Define the extensible attributes. For information, see [Using Extensible Attributes](#).
  7. Click **Next** to schedule a change. In the *Schedule Change* panel, click **Now** to immediately execute this task. Or click **Later** to schedule this task, and then specify a date, time, and time zone.
  8. Click **Save & Close**.



#### Note

After making changes to the extensible attributes, you may need to rebuild the topology EA database. For more information, see [Rebuilding EA Database](#) below.

## Viewing Topology Rulesets

To view topology rulesets, navigate to the **Data Management** tab -> **DNS** tab -> **Traffic Control** tab, and then click **Manage Topology Rulesets** in the Toolbar. The Topology Manager lists the configured rulesets, their destination types, sites, and comments.

You can perform the following:

- Add new rulesets. To add a new ruleset, click the Add icon. For more information, see [Defining Topology Rulesets](#) above.
- To edit a ruleset, click the checkbox next to the ruleset name, and then click the Edit icon. You can modify the following in the *Ruleset* editor:
  - In the **General Basic** tab, you can perform the following:
    - Add new rules to the ruleset. Click the arrow next to the Add icon and select either **Extensible Attribute Rule**, **Subnet Rule**, or **Geography Rule**. For more information, see [Defining Topology Rulesets](#) above.
    - Modify rules in the ruleset. To edit an existing rule, select the checkbox of the required rule in the Rules table, and then click the Edit icon. When you are finished editing, click **Save** above the Rules table. For more information, see [Defining Topology Rulesets](#) above.
    - Delete existing rules from the ruleset. Select the checkbox of the required rule in the Rules table, and then click the Delete icon.  
You can modify the destination type only if there are no rules in the ruleset.
  - In the **Extensible Attributes** tab, you can add new or edit existing extensible attributes. For information, see [Using Extensible Attributes](#).
- Delete a ruleset or schedule the deletion for a later time.

- To delete a ruleset, select the checkbox next to its name and click the arrow next to the Delete icon. To delete the object immediately, select **Delete**.
- To schedule the deletion, click **Schedule Delete**. For more information, see [Scheduling Deletions](#).
- Export topology rulesets. To export the entire list of rulesets in a format that can be imported, click the Export icon and choose **Export data in Infoblox CSV Import format**. To export all data that is currently visible in the Topology Manager, click the Export icon and choose **Export visible data**.
- Print the data that is currently visible in the Topology Manager. Click the Print icon to print.

## Importing a Topology Database

The DNS Traffic Control license includes a MaxMind location database that is deployed when you enable the DNS Traffic Control. Note that only a single MaxMind location database can be present on the Grid at a time. The MaxMind location database contains various geographic locations that can be used when you define a geography rule. NIOS supports both paid GeoIP2 and free GeoLite2 MaxMind location databases. The GeoLite2 MaxMind Country database is shipped with the NIOS appliance. The MaxMind location database is static over the lifetime of the querying process until you import a new database and restart services.

When you import a new MaxMind location database, the appliance replaces the existing database. You can import MaxMind location databases that are in MMDB or CSV format. To view the current version of the database, click **Current Version**.

You can import a ready-to-use MaxMind location database or create your own ZIP file containing multiple CSV files. To import a MaxMind location database or to view the current version of the database, complete the following:

1. From the **Data Management** tab, select the **DNS** tab, and then select the **Traffic Control** tab.
2. Click the arrow next to the **Topology Database**, and then select **Import GeoIP Database** from the drop-down list.
3. In the *Import Topology Database* wizard, complete the following:
  - **File:** Click **Select** and navigate to the MaxMind location database.
  - **Upload:** Click **Upload** to import the MaxMind location database.
4. In the Toolbar, click the arrow next to **Topology Database**, and select **Current Version** from the drop-down list to view the details of the imported MaxMind location databases. In the Geography section, the Grid Manager displays the database type, build date, build version, and the date and time when the database was deployed to the Grid Master.  
The latest database version may not be deployed on all DTC members. To view the current deployed versions, select **Data Management -> DNS -> Members**.

To create a custom database in a ZIP file, complete the following:

1. Create a directory with CSV files and name them using the following pattern:  
{Product}-{Content}-{Blocks-or-Locations}-{version-or-localization}.csv.

Only the three CSV files matching these patterns are suitable for the import:

```
{Product}-{Content}-Blocks-IPv4.csv
{Product}-{Content}-Blocks-IPv6.csv
{Product}-{Content}-Locations-en.csv
```

For example:

```
GeoLite2-City-Blocks-IPv4.csv
GeoLite2-City-Blocks-IPv6.csv
GeoLite2-City-Locations-ru.csv
```

or

```
GeoIP2-Country-Blocks-IPv4.csv
GeoIP2-Country-Blocks-IPv6.csv
GeoIP2-Country-Locations-en.csv
```

where

“GeoLite2” and “GeoIP2” correspond to {Product}

“City” and “Country” correspond to {Content}

“IPv4” and “IPv6” correspond to {version}

“ru” and “en” correspond to {localization}

Note that the locations file and at least one of the Blocks files must exist or the import fails. Also, all of these files

must have identical {Product}-{Content} pairs or the import fails. You can use a ready-to-use MaxMind location database as an example.

2. You can add multiple CSV files for different localizations to your ZIP file. Use the following naming pattern: {Product}-{Content}-Locations-{localization}.csv.  
For example:  
GeoLite2-City-Locations-ru.csv  
GeoIP2-City-Locations-de.csv  
GeoIP2-Country-Locations-en.csv
3. Add the directory with the CSV files to a ZIP file. The name of the ZIP file you upload and the name of the directory in the ZIP file are not significant. The ZIP file should contain only one directory and no subdirectories. Any files in the ZIP file with an extension different from .csv are ignored.
4. Import the ZIP file to Grid Manager as described above.



#### Note

The Country database does not support 'subdivision' labels and importing it invalidates all existing rules that use 'subdivision' labels.

## Rebuilding EA Database

Unlike the GeoIP database, the EA database is not imported externally but configured within the system. After making changes to extensible attributes, Grid Manager offers you to rebuild the DNS Traffic Control Topology Database. You can use the banner that appears at the top of the screen and then click **Rebuild** to rebuild the database immediately. Or, you can click **Ignore** to rebuild the database later in the **Traffic Control** tab. Clicking **Ignore** applies to all changes that require a rebuild of the EA database. The EA database rebuild is ignored for the duration of the user session.

To rebuild the EA database, complete the following:

1. From the **Data Management** tab, select the **DNS** tab, and then select the **Traffic Control** tab.
2. In the Toolbar, click the arrow next to the **Topology Database** and select **Rebuild EA Database -> Rebuild** or **Schedule Rebuild**.
3. In the *Rebuild EA Database* dialog box, select **Yes** to rebuild the database or **No** to discard the rebuild. To schedule the rebuild task, in the *Rebuild EA Database Schedule* dialog box, specify a date, time, and time zone.

To view the current version of the EA database, click **Topology Database -> Current Version** in the Toolbar. Grid Manager displays the database build date and its last rebuild status in the Extensible Attributes section.



#### Note

The latest database version may not be deployed on all DTC members. To view the current deployed versions, select **Data Management -> DNS -> Members**.

## Managing DNS Traffic Control Objects

You can configure multiple DNS Traffic Control servers, pools, or LBDNs on a NIOS appliance. For information on how to configure and manage the DNS Traffic Control Objects, see:

- [Managing DNS Traffic Control Pools](#)
- [Managing DNS Traffic Control LBDNs](#)
- [Managing DNS Traffic Control Servers](#)

Sections covered in the topic are:

- [Viewing DNS Traffic Control Objects](#)
- [Deleting DNS Traffic Control Objects](#)



- [Enabling or Disabling Traffic Control Objects](#)
  - [Disabling Traffic Control Objects](#)
  - [Enabling Traffic Control Objects](#)
- [Limitations and Warnings for Auto Consolidated Monitors](#)
  - [DNS Traffic Control](#)
  - [DNS Traffic Control LBDN](#)
  - [DNS Traffic Control Servers](#)

## Viewing DNS Traffic Control Objects

Grid Manager lists all DNS Traffic Control objects in the **Traffic Control** tab. You can view the objects that you have configured in the Grid from: **Data Management** tab -> **DNS** tab -> **Traffic Control** tab.

Based on the selected columns, Grid Manager displays the following information for each DNS Traffic Control object:


- **Name:** The name of the object.
- **Type:** The object type.
- **Status:** Displays information about the last update, connection status, load balancer methods, and servers and pools. Hover your mouse over the status value to view full information in a tooltip. For more information about the possible statuses, see as described in [Understanding DTC Object Status](#).
- **IPv4 Address:** The IPv4 address of the object, if applicable.
- **IPv6 Address:** The IPv6 address of the object, if applicable.
- **Disabled:** **Yes** or **No**. Indicates whether the DNS Traffic Control object is disabled.
- **Comment:** Displays any comments that were entered for the object.
- **Last Status Update:** Displays the timestamp of the last status update.
- **Load Balancing Method:** Displays the load balancing methods defined for the object.
- **Topology Ruleset:** Displays the topology ruleset defined for the object if it uses the Topology load balancing method.
- Extensible attributes, if configured:
  - **Site:** Displays any values that were entered for the Site pre-defined attribute.
  - **IB Discovery Owned:** Displays any values that were entered for the IB Discovery Owned pre-defined extensible attribute.
  - **Building:** Displays any values that were entered for the Building pre-defined attribute.
  - **Country:** Displays any values that were entered for the Country pre-defined attribute.
  - **Region:** Displays any values that were entered for the Region pre-defined attribute.
  - **State:** Displays any values that were entered for the State pre-defined attribute.
  - **VLAN:** Displays any values that were entered for the VLAN pre-defined attribute.



### Note

You can perform inline editing in the **Name**, **Comment**, and **Site** columns by double-clicking the required line in the table and providing the value in the corresponding column.

You can perform the following in the **Traffic Control** tab:

- Select the checkbox to view specific objects only:
  - **LBDN:** Select the checkbox to view LBDN objects only. For more information, see [Configuring DNS Traffic Control LBDNs](#).
  - **Pool:** Select the checkbox to view pools only. For more information, see [Configuring DNS Traffic Control Pools](#).
  - **Server:** Select the checkbox to view servers only. For more information, see [Configuring DNS Traffic Control Servers](#).
- Change the set of columns displayed in the DTC objects table and change their width. For more information, see [Customizing Tables](#).
- Click the Add icon to add an object.
- Select an object and click the Edit icon to edit the configuration. You can also click the Action icon  of the object and select **Edit** from the menu.



- Select an object and click the Delete icon to delete it. You can also click the Action icon of the object and select **Delete** from the menu. For more information, see the [Deleting DNS Traffic Control Objects](#) section.
- Click a DTC server name to open the list of DTC records associated with the server. For more information, see [Managing DTC Server Records](#).
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Click the Export icon to export the list of objects to a .csv file.
- Click the Print icon to print the list of objects.
- Add or delete extensible attributes for a DTC object by selecting the object in the table and clicking **Extensible Attributes** in the Toolbar. For information, see [Using Extensible Attributes](#).
- Test a selected LBDN by clicking **Test LBDN** in the Toolbar. For more information, see [Testing DNS Traffic Control LBDNs](#).
- Enable/disable one or more selected traffic control objects. For more information, see the [Enabling or Disabling Traffic Control Objects](#) section.
- View a visualization of the traffic control structure for an object by selecting the object in the table. The visualization is displayed by default. To hide the visualization, click **Hide Visualization** in the Toolbar. For more information, see [Visualization for DNS Traffic Control Objects](#).
- Use the **IDN Converter** from the Toolbar to convert IDNs into punycodes. For more information, see [Decoding IDNs and Encoding Punycode](#).



#### Note

The Grid Master Candidate provides the health status of DNS Traffic Control objects such as servers, pools, and LBDNs through WAPI requests.

## Deleting DNS Traffic Control Objects

You can delete DNS Traffic Control objects, such as servers, pools, or LBDNs. When you delete an LBDN, the appliance automatically dissociates it from the zones. To delete an LBDN, you must either have write permission on the LBDN record or the LBDN. For more information, see [License Requirements and Admin Permissions](#).

You cannot delete a DNS Traffic Control pool when it is in use. To delete a pool, you must first delete it from the associated LBDNs. You cannot delete a DNS Traffic Control server when it is in use. You must first remove it from every pool and topology ruleset before deleting the server.

To delete an object:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, click the Action icon of the object name and select **Delete** from the menu or select an object and click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes** to delete the object or **No** to cancel the operation.

The *Delete Confirmation* dialog box displays information about associated objects and warns if the object cannot be deleted due to the associations.

To schedule an object deletion, click the Action icon of the object and select **Delete** -> **Schedule Delete**. Alternatively, you can select the object in the **Traffic Control** panel and then select the Delete icon -> **Schedule Delete**. In the *Schedule Deletion* panel, enter a date, time, and time zone. For information, see [Scheduling Deletions](#).



#### Note

If you remove a name server associated with a zone that comprises LBDN records and if the name server is configured as part of a [consolidated monitor list](#), ensure that you remove the name server from the consolidated health monitor list in the DTC pool. For more information about health monitors, see [Using DNS Traffic Control Health Monitors](#) and [Configuring DTC Monitors for Health Check](#).

## Enabling or Disabling Traffic Control Objects

You can enable or disable a single or multiple DNS Traffic Control objects for a selected Grid member.

The Grid-level status of DNS Traffic Control objects can be viewed on the **Data Management** tab -> **DNS** tab -> **Traffic Control** tab and the member-level status in the visualization panel. For more information, see [Visualization for DNS Traffic Control Objects](#) and [Understanding DTC Object Status](#). You will not be able to view the disabled configuration of an object in Grid Manager.

## Disabling Traffic Control Objects

To disable a traffic control object:



### Note

Ensure that at least one health monitor is configured in a pool or a server object before you attempt to disable that DTC object.

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab.
2. In the *Traffic Control* panel, select the object that you want to disable.
3. Click the arrow beside the Enable/Disable icon or click the Action icon of the selected object, and select **Disable**.
4. Complete the following steps in the *Disable Traffic Management Objects* dialog:
  - **Disable health monitoring:** Select it if you want to disable the health monitoring when the selected DTC object is disabled.
  - **Disable options:** Select one of the following options:
    - **Disable until objects are enabled manually:** Select it to keep the object disabled until you enable the object manually.
    - **Disable until DNS service restarts:** Select it to keep the object disabled until the DNS service is restarted.
    - **Disable for specified time (seconds):** Select it and specify a time interval in seconds until when the object must remain disabled. The object is enabled automatically after the specified interval elapses.
    - **Disable after (seconds):** Select it and specify a time interval in seconds after which the object is disabled automatically.
  - In the **DTC grid members** section, select the Grid members on which the selected object must be disabled:
    - i. Click a member in the **Enabled on** box to select it.
    - ii. Click the right arrow to move the member to the **To disable on** box.
5. Click **Save & Close**.

Consider the following points when you have disabled DNS Traffic Control objects in a Grid:

- When you restore a Grid that had a DNS Traffic Control object disabled with the **Disable until DNS restarts** option, the object gets enabled as the restore operation causes all services including the DNS service to restart. You must disable the object once again.
- If you need to override or modify the configuration used to disable a DTC object, you must either wait until the object is enabled based on the setting you have defined, or manually enable the object.
- Disabling DNS Traffic Control objects does not impact the DTC backup or restore operation. A restore operation enables all objects in a Grid leaving them in Running state.
- When a Dig request is running, if you force restart the DNS service, the disabled DTC objects continue to send responses until all objects that were in the disabled state are reinstated to the disabled state after the restart. The time taken to reinstate the objects varies depending on the count of the disabled objects.
- When you modify the name server association of zones configured for an LBDN object, and then restart the DNS service, the DTC object disable settings (manual fail back events) configured for members associated with the LBDN, its pool, or server objects prior to the zone's name server modification, becomes stale, and the stale data is cleared at regular intervals.

## Enabling Traffic Control Objects

In a DNS Traffic Control setup, you must have at least one server object enabled in a pool and at least one pool object enabled in an LBDN.

To enable a traffic control object:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab.
2. In the *Traffic Control* panel, select the object that you want to enable.
3. Click the arrow beside the Enable/Disable icon, or click the Action icon of the selected object, and select **Enable**.
4. In the *Enable Traffic Management Objects* dialog, select the Grid members on which the object must be enabled:
  - a. In the **DTC grid members** section, click a member in the **Disabled on** box to select it.
  - b. Click the right arrow to move the member to the **To enable on** box.
5. Click **Save & Close**.



### Note

- To enable a disabled DNS Traffic Control server object, its pool must be enabled. If not, the status of the server object will remain disabled.
- In NIOS version 8.6.2 and later, to enable an object that was disabled in a prior version of NIOS, do not use the **Enable** option. Instead, complete the following steps:
  - a. On the **Traffic Control** tab, click the Action icon of the disabled object and select **Edit**.
  - b. In the editor, on the **General** tab, clear the selection in the **Disabled** checkbox.
  - c. Click **Save & Close**.

## Limitations and Warnings for Auto Consolidated Monitors

You can view the limitations and warnings of Auto Consolidated Monitors for the DNS Traffic Control, server, and LBDN objects listed in this section.

### DNS Traffic Control

- The Auto Consolidated Monitors option uses DNS Traffic Control consolidated monitors that cannot share results of health monitors specified on DNS Traffic Control servers (server monitors). Therefore, DNS Traffic Control objects may have some differences in health statuses per Grid members when the Auto Consolidated Monitors option is enabled.
- If you disable the Auto Consolidated Monitors option for a DNS Traffic Control pool that is linked to an LBDN that has Auto Consolidated Monitors enabled, an error message is displayed. You cannot disable the Auto Consolidated Monitors option when this option is enabled in the linked LBDN.
- You cannot make any changes to the DNS Traffic Control consolidated monitors settings if the Auto Consolidated Monitors option is enabled.




### Warning

- If you enable the Auto Consolidated Monitors option on the DNS Traffic Control pool, all existing DNS Traffic Control consolidated monitors are deleted.
- When there are no health monitors on the DNS Traffic Control pool and you enable the Auto Consolidated Monitors option, a warning message asking you to manually add health monitors to the DNS Traffic Control pool is displayed.


### DNS Traffic Control LBDN

- Since the Auto Consolidated Monitors option uses DNS Traffic Control Consolidated Monitors which cannot share results of health monitors specified on DNS Traffic Control Servers (server monitors), DNS Traffic Control objects may have some differences in health statuses per Grid members even when the Auto Consolidated Monitors option is enabled.

 Warning

- If you enable Auto Consolidated Monitors on the LBDN, all the existing DNS Traffic Control Consolidated Monitors on all linked DNS Traffic Control Pools are deleted.
- When Auto Consolidated Monitors is enabled on LBDN, all linked DNS Traffic Control Pools also have this option automatically enabled.
- When there are no health monitors assigned on the linked DNS Traffic Control Pool(s), and you enable Auto Consolidated Monitor on the LBDN, the warning message asking to manually add health monitors on DNS Traffic Control Pool is displayed.

## DNS Traffic Control Servers

 Warning

When you add health monitors to a DNS Traffic Control server that is a part of the DNS Traffic Control configuration with the source IP hash load balancing method selected and the Auto Consolidated Monitors option enabled, a warning message is displayed that these changes can cause differences in resolving queries among DNS Traffic Control Grid members for the same DNS request.

## Managing DNS Traffic Control Pools

A pool contains load balanced servers. You can define multiple servers for a pool. Each LBDN must have at least one pool associated with it to be operational. For sites with a large volume of incoming traffic, you can configure DNS Traffic Control to distribute client requests to multiple servers using a load balancing pool. An individual server can belong to one or multiple load balancing pools, depending on how you want to manage your network traffic. You can also set the order of servers in the pool and define a ratio on a server basis.

## Configuring DNS Traffic Control Pools

A pool can contain preferred and alternative load balancing methods. You can define permissions on these pools and associate extensible attributes with them. Each pool can contain one or more health monitors associated with it. You can define TTLs at the LBDN level. These TTLs are valid for dynamic RRsets that are created by the querying process for each query.

To configure a pool, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, and then click the arrow next to the Add icon and select **Pool**.
2. In the *DTC Pool Wizard*, complete the following:
  - **Name**: Enter the name of the pool.
  - **Comment**: Enter additional information about the pool.
  - **Auto Consolidated Monitors**: Select this option to enable auto-managing of DNS Traffic Control Consolidated Monitors. This option locks all the DNS Traffic Control consolidated monitor settings and creates a DNS consolidated monitor for each health monitor in the pool with the following settings:
    - Availability – ALL
    - Full health communication – Yes
    - Listen To members - all DNS Traffic Control members that serve linked LBDN(s).

Enabling the **Auto Consolidated Monitor** option also helps to synchronize the health statuses of DNS Traffic Control objects in the LBDN for each DNS Traffic Control Grid member, which means if one Grid member considers that the DNS Traffic Control Server is RED, this server will be RED for each Grid member performing the health check. For information on DNS consolidated monitors, see [Configuring DTC Monitors for Health Check](#).

For information on Limitations and Warnings of the **Auto Consolidated Monitors** option, see [Limitations For DNS Traffic Control](#).

- **Disabled**: Select this to disable the pool.
3. Click **Next** to associate health monitors with the pool:

- **Health Monitors:** Select the health monitor from the **Available** table, which you want to associate with the pool, and then click the right arrow to move the selected health monitor to the **Active** table. You can use SHIFT+click and CTRL+click to select multiple health monitors. To dissociate the health monitor from the pool, select it and click the left arrow to move it to the **Available** pane from the **Active** pane.
  - **Availability Requirements:** Select from the following:
    - **All:** All active monitors must report the available status for the pool to be determined as available.
    - **Any:** Any number of active monitors must report the available status for the pool to be determined as available.
    - **At least:** The minimum number of active monitors that must report the available status for the pool to be determined as available.
4. Click **Next** and select the preferred load balancing method:
    - **All Available**
    - **Ratio: Dynamic (see details below in the procedure)**
    - **Global Availability**
    - **Source IP Hash**
    - **Ratio: Fixed**
    - **Round Robin**
    - **Topology (see details below in the procedure)**  
For more information, see [Load Balancing Methods for DNS Traffic Control](#).
  5. If you select **Ratio: Dynamic** as the preferred method, also select a dynamic ratio method from the following:
    - **Round Trip Delay:** Select this to enable load balancing based on the proximity of DTC servers determined through round trip delay. Specify the following:
      - **Monitor:** Select a pre-configured health monitor to use for monitoring the round trip delay.
    - **SNMP:** Select this to enable load balancing based on a server metric captured by an SNMP health monitor. Specify the following:
      - **Monitor:** Select a health monitor for which to track a server metric.
      - **OID:** Specify an object identifier that indicates the metric to track.
      - **Weighing:** Select to weigh DTC servers by priority or ratio.
      - **Inverse OID value:** Select this if you want to use the value of the monitored metric as inverted for convenience of determining the availability of the server.
  6. If you select **Topology** as preferred method, also select a Topology Ruleset. Only topology rulesets with the Server destination type are displayed in the drop-down list.
 

If you select Topology as the preferred method, you can also specify the alternate method which is used to select a server from the pool if the preferred one does not return any result. The preferred and alternate methods must be different.
  7. If applicable, select the alternate load balancing method.
    - **All Available**
    - **Ratio: Dynamic**
    - **Global Availability**
    - **Source IP Hash**
    - **None**
    - **Ratio: Fixed**
    - **Round Robin**
    - **Topology**

For details on each alternate method, see the description of the preferred method above.
  8. Click **Next** to associate servers with the pool. Click the Add icon, select a server from the *DTC Server Selector* dialog box and click **OK**. You can use SHIFT+click and CTRL+click to associate multiple servers. The appliance displays the following information:
    - **Server Name:** The name of the DNS Traffic Control server.
    - **Host:** The host address of the server.
    - **Ratio:** You can modify the ratio value. The value must be greater than zero.
    - **Disabled:** Indicates whether the server is disabled.
    - **Order:** Displays the order of servers in the list.  
To dissociate a server from the pool, select the checkbox next to the server name and click the Delete icon.
  9. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).

10. To schedule the change, click **Next** or **Schedule for Later**. In the *Schedule Change* panel, select **Now** to immediately execute this task. Or, select **Later** to schedule this task, and then specify a date, time, and time zone.
11. Save the configuration.

## Modifying DNS Traffic Control Pools

To modify a pool, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, click the Action icon next to the pool name and select **Edit** from the menu.
2. The *DTC Pool* editor contains the following basic tabs from which you can modify data. For information about how to modify data, see [Configuring DNS Traffic Control Pools](#) above.
  - **General**: This tab displays the pool **Name** and **Comment**. You can edit the values and enable or disable the pool.
  - **TTL**: This tab displays the TTL value configured for the pool. The default value is inherited from the LBDNs which are using the pool. There can be multiple inheritance. Click **Override** to override the value.
  - **Health Monitors**: This tab displays health monitors that are associated with the pool. You can associate new health monitors or dissociate the health monitors that are already associated with the pool.
  - **Load Balancing**: This tab displays the load balancing methods that you have selected while configuring the pool. You can select a new preferred and alternate load balancing methods.
  - **Pool Members**: This tab displays the servers that are associated with the pool. You can add new servers or delete servers that are associated with the pool. You can also modify the ratio and order of servers.
  - **Extensible Attributes**: Add and delete extensible attributes that are associated with the pool. You can also modify the values of extensible attributes. For information, see *Using Extensible Attributes*.
3. To schedule this task, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, select **Later** and enter a date, time, and time zone. The Schedule icon is green when there is a pending scheduled task. For information, see [Scheduling Tasks](#).
4. Save the configuration.

## Managing DNS Traffic Control LBDNs

A DTC LBDN is a load balanced domain name object that is used by DNS Traffic Control to process DNS queries for load balanced resources. You can define multiple LBDNs on the NIOS appliance and associate extensible attributes to them. You can configure permissions for DTC LBDNs. For more information, see [License Requirements and Admin Permissions](#).

## Configuring DNS Traffic Control LBDNs

You can configure a load balancing method for each LBDN. You can assign multiple pools and a single load balancing method to an LBDN. You can associate or dissociate LBDNs with a zone. Note that zone transfers and incremental zone transfers ignore LBDNs. When you configure or modify DTC LBDNs, a service restart is required in order for the new configuration to take effect.

On the appliance, the DNS Traffic Control querying process generates A, AAAA, NAPTR, SRV or CNAME records for an LBDN, called LBDN records. LBDN records are served by DNS Traffic Control servers. An LBDN record must be associated with an authoritative zone. For more information about LBDN records, see [Managing LBDN Records](#) below.

To configure an LBDN, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, click the arrow next to the Add icon, and select **LBDN**.  
or  
From the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *authoritative\_zone* -> **Records** tab, click the Add icon, and select **Record** -> **DTC LBDN**.
2. In the *DTCLBDN* wizard, complete the following:
  - **Display Name**: Enter a display name for the LBDN.
  - **Patterns**: Click the Add icon to add an LBDN pattern. For more information, see [Configuring LBDN Patterns](#).  
To delete an LBDN pattern, select the checkbox next to the pattern and click the Delete icon.



- **Load Balancing Method:** From the drop-down list, select the method you want to use for load balancing. You can select one of the following: **Global Availability**, **Source IP Hash**, **Ratio:Fixed**, **Round Robin**, and **Topology**. The default is **Round Robin**. For more information about the load balancing methods, see [Load Balancing Methods for DNS Traffic Control](#).
- **Topology Ruleset:** This is displayed only when you select the Topology load balancing method. In the drop-down list, only the topology rulesets with the Pool destination type are displayed. Select a topology ruleset for the Topology load balancing method. For more information, see [Defining Topology Rulesets](#).
- **Persistence:** Select this checkbox and enter a value greater than zero seconds to enable persistence for the LBDN. You can specify a period between one second to 2 hours. Even if the DNS restart takes longer than the value specified in the **Persistence** field, the DNS server directs the request to the same server. If you specify zero, the appliance does not cache the requests. When you enable persistence for an LBDN, the appliance stores the results for specific LBDN responses in the DNS Traffic Control cache. When a request originates from the respective FQDN or an IP address within the specified period, the DNS server directs the request to the same server.

Note that when the persistence is enabled, cached results are not guaranteed to persist for the configured duration. The maximum size of the persistence cache is limited globally by the platform. When the limit exceeds the maximum size, the oldest results are deleted. The appliance might discard persistence results if the relevant configuration changes. HA DTC cache replication works on both active and passive nodes and during an HA failover, the DTC cache is replicated from the active node to the passive node. DTC cache replication in HA mode is supported only for IPv4 communications.

- **Priority:** Select a priority value, **1** (High), **2** (Normal), or **3** (Low). The priority value is used when there are LBDNs that have patterns matching the same FQDN and that are assigned to the same zone. In this case, the matching LBDN with the highest priority is used. For example, an LBDN with `*.foo.com` and an LBDN with `www.*.com` patterns can be linked to the same zone `foo.com` if the LBDN with the `*.foo.com` pattern has priority 1 and the LBDN with the `www.*.com` pattern has priority 2 or 3. If there are no matches, the default LBDN is used.
- **Comment:** Enter additional information about the LBDN object.
- **Auto Consolidated Monitors:** Select this option to enable auto managing DNS Traffic Control Consolidated Monitors on DNS Traffic Control Pools linked to the LBDN. This option will lock all the DNS Traffic Control Consolidated Monitor settings on affected DNS Traffic Control Pools, and creates a DNS Traffic Control Consolidated Monitor for each health monitor with the following settings:
  - Availability - ALL
  - Full health communication - Yes
  - Listen To members - all DNS Traffic Control members that serve particular LBDN.

The **Auto Consolidated Monitor** option also helps to synchronize the health statuses of DNS Traffic Control objects in the LBDN for each DNS Traffic Control Grid member, which means if one Grid member considers that the DNS Traffic Control Server is RED, this server will be RED for each Grid member performing the health check.

For information on Limitations and Warnings of the **Auto Consolidated Monitors** Option, see [Limitations for DNS Traffic Control LBDN](#).
- **Disabled:** Select this to disable the LBDN.

3. Click **Next** and complete the following:

- **Return these record types for the associated zones:** Select any or all of the following LBDN record types: **A**, **AAAA**, **NAPTR**, **SRV**, and **CNAME**. You must select at least one record type for the LBDN, otherwise the LBDN is disabled. The patterns and the record types can overlap with another LBDN that is linked to the same zone only if their priorities differ.

If you select the **A** or **AAAA** record type, the LBDN returns the corresponding record and/or a **CNAME** record when the client queries for any record type and if the server selected by DNS Traffic Control has the required data.

However, if the client queries for **CNAME** explicitly, ensure that you select the **CNAME** record type checkbox for the **CNAME** records to be returned.

If you select the **CNAME** or **NAPTR** record type, the LBDN returns the **CNAME** or **NAPTR** record respectively when the client queries for those records and if the server selected by DNS Traffic Control has the required data. As the **CNAME** response must be unique, the **CNAME** record type is unavailable for an LBDN if any pool in that LBDN uses the All Available load balancing method.

Unlike other DNS Traffic Control record types, **SRV** record type has a name. If the QNAME matches the pattern in LBDN and the QTYPE is enabled, a server is selected and all the records of the QTYPE configured for the server are returned. DNS Traffic Control SRV name is not used in name matching during DNS resolution in BIND.

To receive distinct responses, use separate LBDNs as well as separate servers for every service/protocol/domain combination.

- **Associated Zones:** Click the Add icon to associate zones with the LBDN. Select a zone from the *ZoneSelector* dialog box and click **OK**. The appliance displays the following information:
    - **Zones:** The name of the selected zone.
    - **DNS View:** The DNS view associated with the selected zone (if there is more than one DNS view).The LBDN is active only when you associate zones with it. You can associate only authoritative forward-mapping zones with the LBDN. The LBDN must contain at least one matching pattern for the zone. For example, an LBDN with patterns "www.\*.com" and "www.\*.net" may be linked to a zone "foo.com". For more information, see *Managing LBDN Records* below.
  - You can also associate LBDNs with DNSSEC signed zones. For more information, see *Associating LBDNs with DNSSEC Signed Zones* below.
4. Click **Next** and click the Add icon to associate pools with the LBDN. Select a pool from the *DNS Traffic Control Pool Selector* dialog box and click **OK**. The appliance displays the following information:
    - **Name:** The name of the selected pool.
    - **Ratio:** The ratio of the associated server. You can edit this value.
    - **Comment:** Displays information that you specified for the pool.
    - **Members:** Displays the member associated with the pool.
    - **Order:** Displays the order of the pools.To dissociate a pool associated with an LBDN, select the checkbox next to the respective pool name and click the Delete icon.
  5. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).
  6. To schedule the change, click **Next** or **Schedule for Later**. In the *Schedule Change* panel, select **Now** to immediately execute this task. Or select **Later** to schedule this task, and then specify a date, time, and time zone.
  7. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring LBDN Patterns

An LBDN pattern is a domain name. You can define a pattern with multiple \* and ? in any position of the domain name. Note the following about \* and ?:

- A sequence of \*s has the same effect as a single \*.
- A sequence of ?s will match exactly as many octets as there are ?s.
- A \* terminates on label boundaries and will not match a label separator. For example, \*.com matches *foo.com* but not *www.foo.com*.
- A ? does not match a label separator.
- An empty LBDN pattern will match the root and it is automatically changed to "." when you save the LBDN.
- An LBDN pattern matches an FQDN if the entire FQDN matches.
- LBDN patterns may contain special characters. For example, a\032 *b.com* contains two adjacent spaces.
- LBDN patterns do not support IDN and they will not convert Unicode to punycode. You can enter punycode, but note that the LBDN pattern matching does not support punycode.



### Note

There are many cases where the use of wildcards within LBDN patterns is advisable; however, Infoblox recommends using wildcards with caution in the left-most position because it may lead to unexpected behavior or responses. When in doubt, the most predictable behavior comes from using the target domain name as the pattern when configuring the LBDN.



## Managing LBDN Records

In order to manage an LBDN in an authoritative zone, you must enable the authoritative zone and associate it with the LBDN. If an LBDN pattern matches a zone name, the records of type "DTC LBDN Record" are created in that zone as proxies for the LBDN.

To view DTC LBDN records, complete the following:

- Select the **Data Management** tab -> **DNS** tab -> **Zones** tab -> *authoritative\_zone* -> **Records** tab.

The record name is the zone-relative portion of the pattern, including wildcards. For example, when you link an LBDN with patterns "www.\*.com", "www.a\*z.\*.com", "\*.com", "bar\*.net" to zone "foo.com", the appliance creates three LBDN records with names "" (zone origin), "www", and "www.a\*z" in the zone. These records will refer to their respective LBDN. You cannot modify LBDN records. The appliance creates or deletes LBDN records based on the matched LBDN patterns. When you delete an LBDN, the appliance automatically deletes linked zones associated with it and deletes all LBDN records. You can edit the pattern that is associated with the respective LBDN record. For more information, see [Modifying DNS Traffic Control LBDNs](#) below.

Note that an LBDN record is a separate object from the LBDN and each of these have separate permissions. For more information, see [License Requirements and Admin Permissions](#).



### Note

**SRV** record type uses a name. If the QNAME matches the pattern and the QTYPE is enabled, then a server is selected and ALL records of the QTYPE configured for the server are returned. For distinct responses use separate LBDNs as well as separate servers for every service/protocol/domain combination. DTC SRV name is not used in name matching during DNS resolution in BIND.

## Associating LBDNs with DNSSEC Signed Zones

If a zone is DNSSEC signed, you can still associate an LBDN, but some restrictions apply. You can set either Signed or Unsigned mode for the response from DNSSEC signed zones.

The following restrictions apply in the Signed mode:

- You cannot assign an LBDN to a zone or unassign an LBDN from a zone while signing, i.e. key rollover, is in progress for that zone. For information about key rollovers, see [About Key Rollovers](#).
- If an LBDN is assigned to a zone for which signing is in progress, then all changes to that LBDN and its dependent configuration (including pools, servers, and topologies) are prohibited until signing completes. The only thing you can do while signing is in progress is to assign an LBDN already assigned to a signed zone to another unsigned zone.
- An LBDN assigned to a signed zone cannot use the All Available load balancing method or have a pattern with a wildcard in the zone. Also, you cannot sign an unsigned zone with such an LBDN assigned.

In the Unsigned mode, unsigned responses in signed zones are returned.

For more information about how to set the Signed or Unsigned mode, see [Configuring Grid DNS Traffic Control Properties](#).



### Note

You cannot assign any signed zone during staged Grid upgrade if not all of the NIOS appliances have been moved to a new software version. This restriction is working in both Signed and Unsigned modes.

## Modifying DNS Traffic Control LBDNs

To modify an LBDN, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, click the Action icon next to the LBDN name and select **Edit** from the menu.

2. The *DTC LBDN* editor contains the following tabs from which you can modify data. For information about how to modify data, see [Configuring DNS Traffic Control LBDNs](#) above.
  - **General:** This tab displays the **Display Name**, **Patterns**, **Load Balancing Method**, **Persistence**, **Priority**, and **Comment** of the LBDN object. Here you can also disable the LBDN.
  - **Associated Zones and Records:** This tab displays the record types that can be returned for the associated zones, the **TTL**, and the **Associated Zones**. You can select any or all of the following record types: **A**, **AAAA**, **SRV**, and **NAPTR**. Note that the default TTL value is 8 hours and is inherited from the associated zones of the Infoblox Grid. You can override this value or associate new zones with the LBDN to inherit a new value.
  - **Pools:** This tab displays the pools that are associated with the LBDN. You can delete an existing pool or add new pools.
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with the LBDN. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).
3. To schedule this task, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, select **Later** and enter a date, time, and time zone. The Schedule icon is green when there is a pending scheduled task. For information, see [Scheduling Tasks](#).
4. Save the configuration and click **Restart** if it appears at the top of the screen.

### Testing DNS Traffic Control LBDNs

You can select an LBDN and test the DTC response for the respective LBDN. To test an LBDN, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab.
2. In the **Traffic Control** panel, select the LBDN object that you want to test and click **Test LBDN** in the Toolbar.
3. In the *Test DTC LBDN* dialog box, complete the following:
  - **Query Source:** Enter the client IP address.
  - **Query Name:** Enter the FQDN of the record that you are requesting.
  - **Member:** Click **Select** to select the Grid member that will return the response. The dialog box displays the list of members that have a DTC license.
  - **Record Type:** Select a record type, **A**, **AAAA**, **SRV**, or **NAPTR** from the drop-down list.
4. Click **Start**.

The appliance displays the response for the request in the text output area. To clear the response from the text area, click **Clear**.

### Managing DNS Traffic Control Servers

DNS Traffic Control servers are objects that are associated with synthesized A, AAAA, SRV, or CNAME records. DNS Traffic Control servers can be in multiple pools and can be the destination for multiple topology rulesets. You can disable a server while in use, but note that this affects the pools that are associated with the server. You cannot disable a server if it is the last active server for any pool with which it is associated. To disable such a server, first remove it from the associated pools and topology rulesets.

### Configuring DNS Traffic Control Servers

You can add a DNS Traffic Control server on the **Traffic Control** tab. Alternatively, you can do this on the **DNS -> Zones** tab or **Members/Servers** tab by selecting an existing A, AAAA, or host record in the table, and then clicking **Create DTC Server** in the Toolbar or in the record's action menu.

You can also add a DTC server on the **Data Management -> IPAM** tab based on a selected existing A or host record. You can do so on both **IP Map** and **List** subtabs.

If you use multi-tier architecture and want to monitor the availability of separate components of the DNS Traffic Control server, you can add a health monitor for an individual IP address or domain name of the server. You can do it after you have initially configured the server. For information, see [Modifying DNS Traffic Control Servers](#) below.

To configure a DNS Traffic Control server, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab.
2. In the Toolbar, click the arrow next to the Add icon and select **Server**.

3. In the *DTC Server Wizard*, under **Select an existing DNS record or enter the Name and Host fields**, perform one the following:
  - Optionally, click **Select** and choose an existing DNS record which will pre-populate the server information:
    - i. Select a zone using the filter and click **Apply**.
    - ii. Select the record type, **Host**, **A**, or **AAAA**, by which to filter the records list.
    - iii. Click the required record name to select it.
  - Specify the DTC server name and host:
    - **Name**: Enter the name of the DTC server.
    - **Host**: Specify the server host by selecting and specifying one of the following:
      - **IP Address**: The DTC response from the server will contain an auto-created A (IPv4) or AAAA (IPv6) record with this IP address.
      - **Domain Name**: The DTC response from the server will contain an auto-created CNAME record that uses this domain name.  
This step only applies if you create a DTC server from the **Traffic Control** tab. If you create a DTC server on the **DNS -> Zones**, **DNS -> Members/Servers** tab, or **Data Management -> IPAM** tab, the record is already selected so this step is not available in the *DTC Server Wizard*.
4. **Auto-create DTC records**: If this is enabled and the Host field contains an IP address, an A (IPv4), or AAAA (IPv6) record will be created. If the Host field contains a domain name, a CNAME record will be created. If you do not enable auto-created DTC records, you must create those records manually. For more information, see *Managing DTC Server Records* below.  
A record type that corresponds to the Host field must exist in order for the DTC Server to return a response.
5. **Comment**: Enter additional information about the server.
6. **Disabled**: Select this to disable the server.
7. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).
8. To schedule the change, click **Next** or **Schedule for Later**. In the *Schedule Change* panel, select **Now** to immediately execute this task. Or select **Later** to schedule this task, and then specify a date, time, and time zone.
9. Save the configuration.

## Managing DTC Server Records

You can create A, AAAA, CNAME, SRV, and NAPTR records in a DTC server similar to the NAPTR record in a DNS zone.

A NAPTR (Name Authority Pointer) record specifies a rule that uses a substitution expression to rewrite a string into a domain name or URI (Uniform Resource Identifier). A URI is either a URL (Uniform Resource Locator) or URN (Uniform Resource Name) that identifies a resource on the Internet. For information about NAPTR records, see [Managing NAPTR Records](#).

You can assign multiple A, AAAA, and NAPTR records simultaneously to a DTC server or only one CNAME record. An enabled CNAME record cannot coexist with an enabled A, AAAA, or NAPTR record. A disabled CNAME record cannot coexist with an enabled CNAME record.

This section describes how to add, modify, and delete records in a DTC server. It includes the following sections:

- Adding DTC Records
- Viewing DTC Records
- Modifying DTC Records
- Deleting DTC Records

## Adding DTC Records

To add a DTC record, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab.
2. Click the server name to which you want to add the records. Details of the records added to the server are displayed.
3. Click the arrow next to the **Add** icon and select the type of record you want to add. (A Record, AAAA Record, CNAME Record, SRV Record, NAPTR Record, or Unknown record).
4. Enter the following information in the *Add <Record\_Type> Record* wizard:

- For an A record, complete the following:
  - **IP address:** Enter an IPv4 address for the A record on the DTC server.
  - **Comment:** Optionally, enter additional information about the A record.
  - **Disable:** Select this checkbox to disable the record. Clear the checkbox to enable it.
- For an AAAA record, complete the following:
  - **IP address:** Enter an IPv6 address for the AAAA record on the DTC server.
  - **Comment:** Optionally, enter additional information about the AAAA record.
  - **Disable:** Select this checkbox to disable the record. Clear the checkbox to enable it.
- For a CNAME record, complete the following:
  - **Canonical name:** Enter the complete canonical (or official) name of the host.
  - **Comment:** Optionally, enter additional information about the CNAME record.
  - **Disable:** Select this checkbox to disable the record. Clear the checkbox to enable it.
- For a NAPTR record, complete the following:
  - **Service:** Specifies the service and protocol used to reach the domain name that results from applying the regular expression or replacement. You can enter a service or select a service from the list.
  - **Flags:** The flag indicates whether the resulting domain name is the endpoint URI or if it points to another record. Select one of the following:
    - **U:** Indicates that the output maps to a URI.
    - **S:** Indicates that the resulting domain name has at least one SRV record.
    - **A:** Indicates that the resulting domain name has at least one A or AAAA record.
    - **P:** Indicates that this record contains information specific to another application. Leave this blank to indicate that the DNS client must use the resulting domain name to look up other NAPTR records. You can use the NAPTR records as a series of rules that are used to construct a URI or domain name.
  - **Order:** Select an Integer from 10 to 100, or enter a value from 0 to 65535. This value indicates the order in which the NAPTR records must be processed. The record with the lowest value is processed first.
  - **Preference:** Select an Integer from 10 to 100, or enter a value from 0 to 65535. Similar to the **Preference** field in MX records, this value indicates which NAPTR record should be processed first when the records have the same Order value. The record with the lowest value is processed first.
  - **REGEX:** The regular expression that is used to rewrite the original string from the client into a domain name. RFC 2915 specifies the syntax of the regular expression. Note that the appliance validates the regular expression syntax between the first and second delimiter against the Python re module, which is not 100% compatible with POSIX Extended Regular Expression as specified in the RFC. For information about the Python re module, refer to <http://docs.python.org/release/2.5.1/lib/module-re.html>.
  - **Replacement:** This specifies the domain name for the next lookup. The default is a dot (.), which indicates that the regular expression in the **REGEX** field provides the replacement value. Alternatively, you can enter the replacement value in FQDN format.
  - **Comment:** Optionally, enter a descriptive comment for this record.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
- For an SRV Record, complete the following:
  - **Display input as:** Select the format in which you want the SRV record to be displayed. When you select RFC 2782 format, the appliance follows the `_service._protocol.name` format as defined in RFC 2782. When you select Free format, enter the entire name in the Domain field.
  - **Service:** Specify the service that the host provides. You can either select a service from the list or type in a service, if it is not on the list. For example, if you are creating a record for a host that provides FTP service, select `_ftp`. To distinguish the service name labels from the domain name, the service name is prefixed with an underscore. If the name of the service is defined at <http://www.iana.org/assignments/service-names-port-numbers/service-names-port-numbers.xml>, use that name. Otherwise, you can use a locally-defined name.
  - **Protocol:** Specify the protocol that the host uses. You can either select a protocol from the list or type in a protocol, if it is not on the list. For example, if it uses TCP, select `_tcp`. To distinguish the protocol name labels from the domain name, the protocol name is prefixed with an underscore.
  - **Domain:** Specify the name determined by LBDN.
  - **Preview:** After you have entered all the information, this field displays the FQDN, which is the concatenation of the Service, Protocol, and Domain fields.

- **Priority:** Select or enter an integer from 0 to 65535. The priority determines the order in which a client attempts to contact the target host; the domain name host with the lowest number has the highest priority and is queried first. Target hosts with the same priority are attempted in the order defined in the **Weight** field.
  - **Weight:** Select or enter an integer from 0 to 65535. The weight allows you to distribute the load between target hosts. The higher the number, the more that host handles the load (compared to other target hosts). Larger weights give a target host a proportionately higher probability of being selected.
  - **Port:** Specify the appropriate port number for the service running on the target host. You can use standard or nonstandard port numbers, depending on the requirements of your network. You can select a port number from the list or enter an integer from 0 to 65535.
  - **Target:** Enter the canonical domain name of the host (not an alias). For example, [www2.corpxyz.com](http://www2.corpxyz.com)  
In addition, you need to define an A record mapping as the canonical name of the host to its IP address.
  - **Comment:** Enter a descriptive comment for the record.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
  - For an Unknown record:
    - **Domain name:** Enter the complete canonical (or official) name of the host.
    - **Comment:** Optionally, enter additional information about the CNAME record.
    - **Disable:** Select this checkbox to disable the record. Clear the checkbox to enable it.
5. To schedule the change, click **Next** or **Schedule for Later**. In the *Schedule Change* panel, select **Now** to immediately execute this task. Or select **Later** to schedule this task, and then specify a date, time, and time zone.
  6. Save the configuration and click **Restart** if it appears at the top of the screen.

## Viewing DTC Records

To view the records associated with a DTC server, go to the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab -> *DTC server*. The Grid Manager displays the following for each DTC record:

- **Name:** The name of the record.
- **Type:** The type of record.
- **Data:** The data that the record contains. For a NAPTR record, this field displays the following data: Order, Preference, Flags, Services, REGEX, and Replacement.
- **Comment:** Comment that was entered for the record.
- **TTL:** The TTL (time-to-live) value of the record.
- **Disabled:** Indicates if the record is disabled.

You can perform the following:

- Click the Add icon to add a DTC record.
- Select a record and click the Edit icon to edit the configuration. You can also click the Action icon next to the record and select **Edit** from the menu.
- Select a record and click the Delete icon to delete it. You can also click the Action icon next to the record and select **Delete** from the menu.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Click the Export icon to export the list of DTC records to a .csv file.
- Click the Print icon to print the list of DTC records.

## Modifying DTC Records

To modify a DTC record, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab -> *DTC server*.
2. Select the DTC record you want to modify, and click the Edit icon.
3. The *DTC Record* editor contains the following tabs from which you can modify information:
  - **General:** You can modify most of the information, except for the read-only fields, such as the DNS view. For a description of the fields, see Adding DTC Records above.
  - **TTL:** You can modify the TTL setting. For information, see [Specifying Time To Live Settings](#).
4. Save the configuration and click **Restart** if it appears at the top of the screen.



#### Note

When you modify an existing A record, AAAA record, or a Host record connected to a DNS Traffic Control server, new DNS Traffic Control configurations are applied after the DNS Traffic Control health update or after the specified TTL expires for the given record.

## Deleting DTC Records

To delete a DTC record, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab -> *DTC server*.
2. Select the DTC record you want to delete, and click the Delete icon.
3. When the confirmation dialog box displays, select **Yes**.

Grid Manager moves the DTC record to the Recycle Bin, from which you can restore or permanently delete the record. For information, see [Using the Recycle Bin](#).

## Modifying DNS Traffic Control Servers

To modify a DNS Traffic Control server, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, click the Action icon next to the DTC server and select **Edit** from the menu.
2. The *DTC Server* editor contains the following tabs from which you can modify data.
  - **General**: This tab displays the **Name**, **Host**, **Auto-create DTC records**, and **Comment** that you entered while configuring the server. You can enable or disable the server. For information about how to modify the server data, see [Configuring DNS Traffic Control Servers](#). Additionally, specify the Server Name Indication setting:
    - **Use Alternate SNI Hostname**: If the SNI name is different from what is configured in the address field, select this checkbox and enter the required hostname by which an HTTPS health monitor should connect to the server.
    - **Health Monitors**: Define health monitors for the DTC server:
      1. Click the Add icon. A new row appears in the table.
      2. In **Health Monitor**, select the monitor type: icmp, http, https, sip, pdp, or snmp.
      3. In **Domain Name or IP Address**, type either the FQDN or the IP address to monitor.
      4. If required, add more health monitors for the server as described above. You can add up to ten health monitors per server.  
In **Health Monitors from Pools**, you can see other health monitors assigned to the pools that the server belongs to. The availability requirement for the pools must be set to either "All" or "Any" for you to be able to add server-specific health monitors. For information, see [Configuring DNS Traffic Control Pools](#).
      5. Save the configuration.
    - **Extensible Attributes**: Add and delete extensible attributes that are associated with the server. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#). For information on warnings related to **Auto Consolidated Monitor**, see [Warnings for DNS Traffic Control Servers](#).
3. To schedule this task, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, select **Later** and enter a date, time, and time zone. The Schedule icon is green when there is a pending scheduled task. For information, see [Scheduling Tasks](#).
4. Save the configuration.

## Visualization for DNS Traffic Control Objects

Grid Manager provides a visual tree view that you can use to quickly understand an overall traffic control structure of a selected DTC object. The visualization panel is displayed by default on the **Traffic Control** tab.

To view a visualization of the DNS traffic control structure for an object:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab.



2. Select the checkbox next to the DTC object for which you want to view, for example, LBDN. The LBDN structure tree is displayed in the visualization panel.

The tree shows the selected object and its relationships with other associated DTC objects. The DTC objects are represented as nodes in the hierarchical order of LBDN -> pool -> server levels. Note that when a DTC server is associated with multiple pools, the appliance displays the connection to the selected pool only, not showing the other pools that the server is assigned to. If an LBDN has more than one pool associated, it displays the servers for only one pool at a time. Clicking on another pool of the LBDN displays the servers for that pool.

You can hover your mouse over an object to display a tooltip that contains the following information:

- The name and status of the DTC object. For more information about the possible statuses, see [Understanding DTC Object Status](#) below.
- **Load Balancing Method:** This is displayed only for LBDN objects.
- **Preferred Load Balancing Method:** This is displayed only for pool objects.
- **Alternate Load Balancing Method:** This is displayed only for pool objects.
- **Host:** The domain name or IP address of a server object.
- **Last Updated:** The timestamp when the object was last discovered.
- **Health Monitors:** This displays health monitors status for pools and servers. The health status displayed for a DTC server includes health monitors assigned through the pool as well as individual health monitors per IP address or domain name of the server, if assigned. For example, `monitor_name(monitor_type)` for a pool monitor, and `IP_address_or_FQDN(monitor_type)` for an individual server monitor.
- **Ratio: Dynamic:** In DTC pools and servers visualization, this displays traffic distribution across members and servers with the Ratio: Dynamic load balancing method applied.
- **Number of Pools:** Shows the number of pools for the LBDN.
- **Number of Servers:** Shows the number of servers for the pool.
- **Used by these LBDNs:** Shows all LBDNs that use the pool.
- **Used by these Pools:** Shows all pools that use the server.

To hide or show the visualization panel, click **Hide Visualization** or **Show Visualization** in the Toolbar.

From a DTC object tooltip, you can perform certain actions on the object or add the default visualization in the panel as discussed later in this topic



#### Note







Grid Manager can display a maximum of 100 nodes for each level associated with the currently visualized node.

## Understanding DTC Object Status

The DTC object status can be one of the following:

Icon (Visualization Panel)	Status and Meaning (Visualization Panel)	Status and Meaning (Traffic Control Tab)
	<b>Running:</b> The object is fully available and operational.	<b>Running</b> with a green background: The object is fully available and operational.
	<b>Warning:</b> The object has a warning message. You can check the syslog for any messages.	<b>Warning</b> with a yellow background: The object has a warning message. You can check the syslog for any messages.
	<b>Error:</b> The object has an error. You can check the syslog for any messages.	<b>Error</b> with a red background: The object has an error. You can check the syslog for any messages.



Icon (Visualization Panel)	Status and Meaning (Visualization Panel)	Status and Meaning (Traffic Control Tab)
	<b>Disabled</b> (in the legend) and <b>Requires Manual Enabling</b> (on mouse hover): The object is disabled due to a configuration setting and it must be enabled manually.	<b>Requires Manual Enabling</b> with a white background: The object is disabled due to a configuration setting and it must be enabled manually.
	<b>Temporarily Disabled</b> (in the legend) <b>Will be enabled at &lt;time_stamp&gt;</b> (On mouse hover): The object is disabled for a specified duration due to a configuration setting. It will be automatically enabled at the displayed time.  Or <b>Temporarily Disabled</b> (On mouse hover): The object is disabled due to a configuration setting and will be enabled when the DNS service restarts.	<b>Temporarily Disabled</b> with a dark grey background: The object is disabled due to a configuration setting. This might be due to different reasons, such as the "Disable" flag being set, the DNS service not running on the selected member, a zone not assigned to an LBDN, or the LBDN not associated with a zone for the selected DTC member.  When multiple DTC Grid members are configured with different disable options for a DTC object, collectively the status shows as <b>Temporarily Disabled</b> . You can view the member level status in the visualization panel.
	<b>Disabled, Working</b> (in the legend) and <b>Will be disabled at &lt;time_stamp&gt;</b> (On mouse hover): The object is working fine, but it will be disabled at the displayed time.	<b>Running</b> with a green background: The object is fully available and operational.
	<b>Disabled, Error</b> (in the legend) and <b>Will be disabled at &lt;time_stamp&gt;</b> (On mouse hover): The object has failed. It will be disabled at the displayed time.	<b>Error</b> with a red background: The object has an error. You can check the syslog for any messages.
	<b>Unknown:</b> The DTC object status has not yet been determined.	<b>Unknown:</b> The DTC object status has not yet been determined.
	<b>Unlicensed:</b> The object does not have a DNS Traffic Control license.	<b>Unlicensed:</b> The object does not have a DNS Traffic Control license.

It may take a few minutes for the status of an object to be updated after a configuration change. Grid Manager calculates the status differently for the DTC objects list view and the visualization.

## Calculating the Status

A DTC pool may use multiple health monitors and it has an availability option for these monitors that is used to determine the status of a DTC server. For example, if the availability requirement is "All", then the server is considered to be 'Running' if all the pool health monitors report the server as 'Running'. Servers may appear in multiple pools that have different health monitors and different availability settings. Thus, a DTC server status can be different for different pools. You can filter the DTC objects visualization by all members or a specific member. A DTC Grid member uses health monitors to determine the status of a DTC server. A Grid member is associated with an LBDN if that Grid member is assigned a name server for a zone that is associated with an LBDN. A DTC server is associated with an LBDN if it is part of a DTC pool which is used by the LBDN. Thus, if a DTC Grid member is associated to a DTC server through an LBDN, the Grid member checks the status of the DTC server using the health monitors associated with the DTC server's pool. When the visualization is filtered by a specific DTC Grid member, the displayed status of the DTC servers is the status that was determined by that specific DTC Grid member. When you select **All Members** in the visualization, the displayed status is an aggregated status across all DTC Grid members. For the DTC objects list view, the status is calculated as follows:

- Server: Aggregated status from all pool monitoring across all DTC Grid members.
- Pool: Aggregated status for all its servers across all DTC Grid members.
- LBDN: Aggregated status for all its pools across all DTC Grid members.

For the visualization panel, the status is calculated as follows:

- Server: Status from selected pool monitoring for selected DTC Grid member.
- Pool: Aggregated status for all its servers for selected DTC Grid member.
- LBDN: Aggregated status for all its pools for selected DTC Grid member.

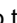


#### Note

If a pool has no health monitors configured, then the servers and pool report a status of 'Running'.

## Working with the Visualization Panel

You can do the following in the visualization panel:

- Resize the visualization panel in relation to the DTC objects list view by dragging the vertical separator between the panel and the list view.
- Filter the visualization by Grid members associated to DTC objects through LBDNs: In the **Member** drop-down list, select **All Members** or a specific member. For more information about how the object status is determined for all members and for a specific member, see Calculating the Status above. Note that the appliance displays only members that have the DNS Traffic Control license.
- Filter by the DTC status: **All DTC Status** or **Non-Running** only. You can combine this filter with the members filter.
- Zoom the map in and out by moving the Zoom slider. You can also zoom in and out by using the mouse wheel.
- Click the **Adjust tree size to window size** icon to adjust the tree size after you zoomed it in or out.
- Change the tree orientation by clicking the Change Tree Orientation icon. The default orientation is vertical.
- Click the **Expand Visualization** icon to open the DTC structure in a separate window. Alternatively, you can click the Action icon  next to the required DTC object in the table and select **Expand Visualization**.
- Click **Show Legend/Hide Legend** to show or hide the legend when the visualization is expanded.
- Click the Refresh icon to refresh the tree. You can also select the **Auto Refresh** checkbox to turn on auto-refresh.
- Click anywhere in the tree and hold your mouse to drag the tree to a desired location in the panel or window.
- Hover your mouse over an LBDN to display the tooltip and do the following:
  - Test the LBDN.
  - Add an existing pool to the LBDN.
  - Add a new pool to the LBDN.
  - Disable or enable the LBDN.
  - Edit the LBDN.
  - Delete the LBDN.
  - Schedule the deletion of the LBDN.
  - Switch to the LBDN visualization mode if you are currently in the pool visualization mode.
- Hover your mouse over a pool to display the tooltip and do the following:
  - Add the pool to an LBDN.
  - Add an existing server to the pool.
  - Add a new server to the pool.
  - Disable or enable the pool.
  - Edit the pool.
  - Switch to the pool visualization mode if you are currently in the LBDN or server visualization mode.
- Hover your mouse over a server to display the tooltip and do the following:
  - Add the server to a pool.
  - Disable or enable the server.
  - Edit the server.
  - Switch to the server visualization mode if you are currently in the LBDN visualization mode.

## Adding Default Visualization

The default visualization allows you to design a DNS Traffic Control structure in the inverse order—first, add a visualization of the default DTC objects structure and create the default disabled objects, and then define the objects one by one.

To add the default visualization:

1. On the **Traffic Control** tab, click the arrow next to the Add icon and select **Default Visualization**.  
The default basic DTC objects structure is displayed in the visualization panel. It consists of the default server, pool, and LBDN. By default, they are disabled.  
The corresponding server, pool, and LBDN objects are added in the DTC objects list view.
2. Configure the DTC objects in any of the following ways:  
For information about the configurable properties of DTC objects, see [Managing DNS Traffic Control Objects](#).
  - Hover your mouse over an object to display the tooltip, click the required button, and make the necessary configurations.
  - In the DTC objects list view, select the required object, click the Edit icon, and make the necessary configurations.
3. If necessary, configure the topology rulesets, topology database, and DTC health monitors used in your DTC structure. For information, see [Defining Topology Rulesets](#), [Importing a Topology Database](#), and [Using DNS Traffic Control Health Monitors](#) correspondingly.
4. After you configured the DTC objects, enable each one of them:
  - Hover your mouse over an object.
  - Click **Enable** in the tooltip.

Once all objects are enabled and all necessary service restarts are performed, the whole DTC structure starts working.

## Backing Up DTC Configuration Files

Infoblox recommends that you regularly back up DTC configuration files. You can either schedule a backup to run at a designated date and time or you can manually back up the files. You can back up the DTC files to the following:

- A local directory
- A TFTP server
- An FTP server. This option requires that you have a valid username and password on the server prior to backing up files.
- An SSH server that supports SCP. This option requires that you have a valid username and password on the server prior to backing up files.

For information about these options, see [Backing Up and Restoring Configuration Files](#).

## Automatically Backing Up DTC Files

To automatically back up DTC configuration files:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Backup -> DTC Backup -> Schedule Backup** from the Toolbar.
2. In the *Schedule DTC Backup* dialog box, select the destination of the backup file from the **Backup to** drop-down list:
  - **TFTP**: Back up system files to a TFTP server.
    - **Keep local copy**: Select this to also save a local copy of the backup file on your appliance. This is disabled by default. The local backup contains only the Grid backup, it does not contain backups for reporting and Network Automation. Note that when you select this, the total backup time will increase.
    - **IP Address of TFTP Server**: Enter the IP address of the TFTP server to which you want to back up the system files.
    - **Directory Path**: Enter the directory path of the file. For example, you can enter `/archive/backups`. The directory path cannot contain spaces. The folder or directory you enter here must already exist on the specified server. Do not include the file name in the directory path.

- **Recurrence:** Select how often you want to back up the files. You can select **Weekly**, **Daily**, or **Hourly** from the drop-down list. When you select **Weekly**, complete the following:
  - **Every:** Choose a day of the week from the drop-down list.
  - **Time:** Enter a time in the hh:mm:ss AM/PM format. You can also click the clock icon and select a time from the drop-down list. The Grid Master creates a backup file on the selected day and time every week.  
When you select **Daily**, enter a time in the hh:mm:ss AM/PM format. You can also select a time from the drop-down list.  
When you select **Hourly**, complete the following:
    - **Minutes after the Hour:** Enter the minute after the hour when the Grid Master creates a backup file. For example, enter 5 if you want the Grid Master to create a backup file five minutes after the hour every hour.
- **Disable Scheduled Backup:** Select this if you want to disable automatic backups from occurring now. You can still save the settings for future use.
- **FTP:** Back up system files to an FTP server.
  - **Keep local copy:** Select this to also save a local copy of the backup file on your appliance. This is disabled by default. The local backup contains only the Grid backup, it does not contain backups for reporting and Network Automation. Note that when you select this, the total backup time will increase.
  - **IP Address of FTP Server:** The IP address of the FTP server.
  - **Directory Path:** Enter the directory path of the file. For example, you can enter **/archive/backups**. The directory path cannot contain spaces. The folder or directory you enter here must already exist on the specified server. Do not include the file name in the directory path.
  - **Username:** Enter the username of your FTP account.  
Note if you have configured AD server for authentication, you must specify "domain name\username".
  - **Password:** Enter the password of your FTP account.
  - **Recurrence:** Select how often the scheduled backups should occur. You can select **Weekly**, **Daily**, or **Hourly**.
  - **Every:** Choose a day of the week from the drop-down list.
  - **Disable Scheduled Backup:** Select this if you want to disable automatic backups from occurring now, but want to save the settings for future use.
- **SCP:** Back up system files to an SSH server that supports SCP.
  - **Keep local copy:** Select this to also save a local copy of the backup file on your appliance. This is disabled by default. The local backup contains only the Grid backup, it does not contain backups for reporting and Network Automation. Note that when you select this, the total backup time will increase.
  - **IP Address of SCP Server:** The IP address of the SCP server.
  - **Directory Path:** Enter the directory path of the file. For example, you can enter **/archive/backups**. The directory path cannot contain spaces. The folder or directory you enter here must already exist on the specified server. Do not include the file name in the directory path.
  - **Username:** Enter the username of your SCP account.  
Note if you have configured AD server for authentication, you must specify "domain name\username".
  - **Password:** Enter the password of your SCP account.
  - **Recurrence:** Select how often the scheduled backups should occur. You can select **Weekly**, **Daily**, or **Hourly**. For information, see the TFTP section.
  - **Every:** Choose a day of the week from the drop-down list.
  - **Disable Scheduled Backup:** Select this if you want to disable automatic backups from occurring now. You can still save the settings for future use.  
When you select **FTP** or **SCP**, ensure that you have a valid user name and password on the server prior to backing up the files.
- **Grid Master (local):** Back up to a local directory on the Grid Master. This is the default.  
By default, the Grid Master generates a backup file and saves it locally in its own storage at 3:00 AM daily. Be aware that backing up the Grid and saving it locally on an hourly basis increases the turnover of files stored on the Grid Master. Backing it up hourly to a remote server increases the overall amount of traffic on your network.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Manually Backing Up DTC Files

You can manually back up DTC configuration files in addition to scheduling your backups. To back up manually:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Backup -> DTC Backup -> Manual Backup** from the Toolbar.
2. In the *Manual DTC Backup* wizard, select the destination of the backup file from the **Backup to** drop-down list:
  - **My Computer**: Back up system files to a local directory on your computer. This is the default.
  - **TFTP**: Back up system files to a TFTP server.
    - **Filename**: Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30`.
    - **IP Address of TFTP Server**: Enter the IP address of the TFTP server to which you want to back up the system files.
  - **FTP**: Back up system files to an FTP server.
    - **Filename**: Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30`.
    - **IP Address of FTP Server**: The IP address of the FTP server.
    - **Username**: Enter the username of your FTP account.  
Note if you have configured AD server for authentication, you must specify "domain name//username".
    - **Password**: Enter the password of your FTP account.
  - **SCP**: Back up system files to an SSH server that supports SCP.
    - **Filename**: Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30`.
    - **IP Address of SCP Server**: The IP address of the SCP server.
    - **Username**: Enter the username of your SCP account.  
Note if you have configured AD server for authentication, you must specify "domain name//username".
    - **Password**: Enter the password of your SCP account.  
When you select **FTP** or **SCP**, ensure that you have a valid username and password on the server prior to backing up the files.
3. Click **Backup**.

## Downloading DTC Backup Files

You can save an existing backup file, or create and save a new one to your local management system, a TFTP server, an FTP server, or a SCP server.

To download an existing backup file:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Backup -> DTC Backup -> Manage Local Backup** from the Toolbar. Grid Manager displays the current backup files in the *Manage Local Backups* dialog box.
2. To download a backup file, select the checkbox of a backup file, and then click the Transfer icon. You cannot select multiple files for downloading.
3. Select one of the following from the **Backup to** drop-down list:
  - **My Computer**: Backup to a local directory on your computer. This is the default.
  - **TFTP**: Save the backup file to a TFTP server.
    - **Filename**: Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30`.
    - **IP Address of TFTP Server**: Enter the IP address of the TFTP server to which you want to save the backup file.
  - **FTP**: Save the backup file to an FTP server.
    - **Filename**: Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30`.
    - **IP Address of FTP Server**: The IP address of the FTP server.
    - **Username**: Enter the username of your FTP server account.
    - **Password**: Enter the password of your FTP server account.
  - **SCP**: Save the backup file to an SSH server that supports SCP.

- **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30`.
  - **IP Address of SCP Server:** The IP address of the SCP server.
  - **Username:** Enter the username of your SCP server account.
  - **Password:** Enter the password of your SCP server account.
- When you select **FTP** or **SCP**, ensure that you have a valid username and password on the server prior to backing up the files.

4. Click **Transfer Copy**.

## Restoring DTC Backup Files

You must log in with a superuser account to back up and restore files. NIOS provides three ways to restore a backup file:

- From a local directory or the management system you use to operate the appliance
- From a TFTP server
- From a remote server using FTP. This option requires that you have a valid username and password on the FTP server prior to performing a backup or restore.



### Note

When you restore NIC interfaces to a VM, ensure that you provision appropriate NIC interfaces with the database content that must be restored to avoid any errors.

To restore a DTC backup file to the same independent appliance or Grid Master:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Restore** -> **Restore DTC** from the Toolbar.
2. In the *Restore* dialog box, choose one of the following from the **Restore from** drop-down list:
  - **My Computer:** Restore a file from your local computer. This is the default.
    - **Filename:** Click **Select File** to navigate to the configuration file.
  - **TFTP:** Restore a file from a TFTP server.
    - **Filename:** Enter the directory path and the file name you want to restore. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30`.
    - **IP Address of TFTP Server:** Enter the IP address of the TFTP server from which you restore the configuration file.
  - **FTP:** Restore a file from an FTP server.
    - **Filename:** Enter the directory path and the file name of the backup file. For example, you can enter `/archive/backups/Infoblox_2009_10_20_15_30`.
    - **IP Address of FTP Server:** The IP address of the FTP server.
    - **Username:** Enter the username of your FTP server account.
    - **Password:** Enter the password of your FTP server account.
  - **Grid Master (Local):** Restore from a local directory on the Grid Master. In the *Backup Set* table, select the file you want to restore.
3. To restore NIOS configuration data, select the **NIOS data** checkbox.
4. To download a backup file from one appliance to a different appliance, select **Force Restore from Different Grid** to enable the feature, and then select one of the following:
  - **Retain Current Grid Master IP Settings** (this is the default)
  - **Overwrite Grid Master IP Settings**
5. Click **Restore**. In the *Confirm Restore* dialog box, click **Yes**.  
After restoring the file, the appliance restarts. The restore process overwrites all existing data. All pending scheduled tasks are not restored or reverted.
6. Close your current browser window, wait a few minutes, and then reconnect to the NIOS appliance.

## Downloading DTC Backup Files from a Different Appliance

When you "force restore" a NIOS appliance, you download a backup file from one appliance to a different appliance. To restore a backup file to the same appliance or Grid Master, use the Restore function as described in the previous section,

see Restoring DTC Backup Files.

To download a backup file from one appliance to a different appliance:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Restore** -> **Restore DTC** from the Toolbar.
2. In the *Restore* wizard, do the following:
  - **Restore from:** Choose a source from which you restore the configuration file, as described in as described in the previous section, see Restoring DTC Backup Files.
3. Select **Force Restore from Different Grid** to enable the feature, and then select one of the following:
  - **Retain Current Grid Master IP Settings** (this is the default)
  - **Overwrite Grid Master IP Settings**
4. Click **Restore**. In the *Confirm Restore* dialog box, click **Yes**.  
After restoring the file, the appliance reboots. The restore process overwrites all existing data. All pending scheduled tasks are not restored or reverted.
5. Close your current browser window, wait a few minutes, and then reconnect to the NIOS appliance.

## Using DNS Traffic Control Health Monitors

Health monitors determine the availability of DTC servers. You associate health monitors with pools. Every health monitor checks each server that is associated with the pool. Additionally, if you use multi-tier architecture, you can assign health monitors to individual IP addresses of a DTC server. You can use pre-defined health monitors or create custom monitors. You can configure health monitors of the following types:

- HTTP health monitors
- ICMP health monitors
- PDP health monitors
- SIP health monitors
- SNMP health monitors
- TCP health monitors

For more information, see the following sections:

- [Configuring HTTP Health Monitors](#)
- [Configuring ICMP Health Monitors](#)
- [Configuring PDP Health Monitors](#)
- [Configuring SIP Health Monitors](#)
- [Configuring SNMP Health Monitors](#)
- [Configuring TCP Health Monitors](#)

When you install the DNS Traffic Control license, each Grid member that is associated with an LBDN independently monitors the health of a server. The appliance resolves the server FQDN using the system resolver. The appliance performs a health check on both the IPv4 and IPv6 addresses. If the member does not have the same IP address types as the server or if the DNS name resolution for A or AAAA records fail to return results, the corresponding health check is considered to be a failure. The appliance caches the addresses that are resolved.

Note that monitoring is done by each Grid member that has a DNS Traffic Control license and is associated with a zone that has an LBDN record. This implies that firewall policies should be such that the member can reach every server it is monitoring. Otherwise, DNS Traffic Control cannot direct responses to servers that are not accessible.

By default, all monitor checks are initiated from the virtual interface of the member that is performing the health check. If the monitored server has only an IPv4 or IPv6 address, then the interface must have the corresponding IP address type. You can configure the DTC monitor source to use the VIP, MGMT, LAN2 (WHERE), or ANY (normal routing) NIOS network interface, or one of the loopback additional IP addresses for the Grid member. You can configure the loopback IP address in the Grid member editor or network configuration.

The status of a DTC server for a specific pool depends upon the status of all the health monitors that are checking it. The status of a pool depends upon the status of all the servers in the pool. The status of an LBDN depends upon the status of all the pools assigned to the LBDN.

For the HTTPS and SIP monitor types, you can upload client certificates and associate them with the monitors to provide when connecting to a DTC server. For information, see [ManagingHealthMonitorCertificates](#) below.

Vice versa, DTC servers provide certificates to authenticate themselves to the HTTPS and SIP monitors. You enable DTC server certificate validation when configuring HTTPS and SIP health monitors.



## Configuring HTTP Health Monitors

An HTTP health monitor sends either an HTTP or HTTPS request to the server. The health monitor then examines the response received from the server. The validation is successful if the server returns a response with the expected result code.

The HTTP/HTTPS monitor can validate the response code and response content. The response content is checked only when the response code is valid. You can define regular expressions to use for the response content check. The supported regular expression syntax is POSIX Extended Regular Expression. For information, see [Supported Expressions for Search Parameters](#).

If the DTC server certificate validation is enabled in the HTTPS health monitor, you can use the Server Name Indication (SNI) feature for remote DTC servers. SNI is an extension to the TLS computer networking protocol by which a client indicates which hostname it attempts to connect to at the start of the handshaking process. This allows a server to present multiple certificates on the same IP address and TCP port number. Thus, multiple secure (HTTPS) websites (or any other service over TLS) can be served off the same IP address without all those sites having to use the same certificate.

After you configure an HTTP/S monitor, you can test its performance. See [Testing HTTP Health Monitors](#) below.



### Note

The HTTP health monitor does not support user name or password authentication.

To configure an HTTP health monitor, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, and then click **Manage Health Monitors** in the Toolbar.
2. In the Health Monitors Manager, click the arrow next to the Add icon and select **HTTP Health Monitor**.
3. In the *HTTP Health Monitor Wizard*, complete the following:
  - **Name:** Enter a name for the HTTP monitor.
  - **Interval (seconds):** Enter the interval value in seconds. The interval value is measured from the end of the previous monitor cycle. The default value is five.
  - **Timeout (seconds):** Enter the timeout value in seconds. The monitor waits for the number of seconds that you specify after sending a response. If it does not receive a response within the number of seconds that you specify, then the appliance considers this check as failed. The monitor discards any response it receives after the timeout. The default value is 15.
  - **Retry Up Count:** Enter a retry up count integer value. Retry up count is a value that determines how many valid responses or good health checks in a row must be received by the Grid member from the DTC server for setting the DTC server health status to green. When you specify a value, the appliance computes the duration for which health check must be performed based on the following:  
`interval*retry up count`  
For example, If the DTC server has had a red status for a long time because all the HTTP health checks have failed, and when the Grid receives the first good HTTP health check result, the health check counter is set as 1; however, the status of the DTC server is still retained as red. If you specify the **Retry Up Count** as 3 and the health check interval as 5, when three consecutive good health check results are received, the health check counter value becomes 3 and the DTC server health status now changes to green (3 good health check results in a row at an interval of 5 seconds each over a span of 3\*5=15 seconds).
  - **Retry Down Count:** Enter a retry down count integer value. It is the opposite of the retry up count. Retry down count is a value that determines how many red status health checks (server is unavailable) in a row must be collected by the Grid member from the DTC server to switch the health status from green to red. For example, for a Grid member whose connection to the DTC server is not stable and has its Retry Down Count set as 5, if the health check results received by the Grid member is in the sequence: green, red, red, red, green, then the health status of the DTC server will be retained as green.  
Note that red health status is set when the health monitor reaches the timeout value that is `[health check interval + timeout]` seconds without a valid response.
  - **Comment:** Enter information about the HTTP monitor.

4. Click **Next** and complete the following:

**Port:** For HTTP, the appliance displays port number 80 by default. If you select Use HTTPS, the appliance displays 443 by default.

**Use HTTPS:** Select the checkbox to enable HTTPS. Specify any of the following options that become available when you select to use HTTPS:

- **Client Certificate:** Optionally, you can select a certificate to use while opening the SSL connection for HTTPS. The monitor does not inspect or validate the server certificate, if any. For information about how to upload certificates, see [ManagingHealthMonitorCertificates](#) below.
- **Ciphers:** Specify a list of SSL ciphers in an OpenSSL format. You can specify cipher texts up to 1024 characters. The client certificate and cipher list are only used for HTTPS transport.

The following example commands list some available ciphers:

Example 1:

```
$ openssl ciphers 'HIGH:!DES'  
DHE-RSA-AES256-SHA:DHE-DSS-AES256-SHA:DHE-RSA-CAMELLIA256-SHA:  
DHE-DSS-CAMELLIA256-SHA:ADH-AES256-SHA:ADH-CAMELLIA256-SHA:AES256-  
SHA: CAMELLIA256-SHA:PSK-AES256-CBC-SHA:EDH-RSA-DES-CBC3-SHA:  
EDH-DSS-DES-CBC3-SHA:ADH-DES-CBC3-SHA:DES-CBC3-SHA:DES-CBC3-MD5:  
PSK-3DES-EDE-CBC-SHA:KRB5-DES-CBC3-SHA:KRB5-DES-CBC3-MD5:DHE-RSA-  
AES128-SHA: DHE-DSS-AES128-SHA:DHE-RSA-CAMELLIA128-SHA:DHE-DSS-  
CAMELLIA128-SHA:  
ADH-AES128-SHA:ADH-CAMELLIA128-SHA:AES128-SHA:CAMELLIA128-SHA: PSK-  
AES128-CBC-SHA
```

Example 2:

```
$ openssl ciphers 'DEFAULT:!EDH+aRSA'  
DHE-DSS-AES256-SHA:DHE-DSS-CAMELLIA256-SHA:AES256-SHA:CAMELLIA256-  
SHA:  
PSK-AES256-CBC-SHA:EDH-DSS-DES-CBC3-SHA:DES-CBC3-SHA:PSK-3DES-EDE-  
CBC-SHA: KRB5-DES-CBC3-SHA:KRB5-DES-CBC3-MD5:DHE-DSS-AES128-SHA:DHE-  
DSS-SEED-SHA: DHE-DSS-CAMELLIA128-SHA:AES128-SHA:SEED-SHA:CAMELLIA128-  
SHA:  
PSK-AES128-CBC-SHA:RC4-SHA:RC4-MD5:PSK-RC4-SHA:KRB5-RC4-SHA:KRB5-RC4-  
MD5: EDH-DSS-DES-CBC-SHA:DES-CBC-SHA:KRB5-DES-CBC-SHA:KRB5-DES-CBC-  
MD5:  
EXP-EDH-DSS-DES-CBC-SHA:EXP-DES-CBC-SHA:EXP-RC2-CBC-MD5:EXP-KRB5-RC2-  
CBC-SHA: EXP-KRB5-DES-CBC-SHA:EXP-KRB5-RC2-CBC-MD5:EXP-KRB5-DES-CBC-  
MD5:EXP-RC4-MD5: EXP-KRB5-RC4-SHA:EXP-KRB5-RC4-MD5
```

Note:

The DHE cipher list family ("Diffie-Hellman key agreement" plus "RSA authentication") could consume excessive CPU and is excluded from the defaults used by DNS Traffic Control health monitors. Although you can enable these ciphers by explicitly configuring them in the cipher list for HTTPS and SIP monitors, you should be aware that doing so will increase CPU usage. Since health monitoring in general does not require high security, Infoblox recommends that you enable these ciphers only for target servers that do not accept other types of ciphers.

- **Enable Certificate Validation:** It is highly recommended to select this for the DTC server certificate to be validated by NIOS.
    - **Enable SNI (Server Name Indication):** Specify if you want to use SNI for the health monitor to connect to a specific DTC server by hostname. In addition, you should indicate an alternate SNI hostname in the DTC server editor.
5. Click **Next** and complete the following:
- **HTTP Request:** Specify the HTTP request to send the query from the client to the server. The appliance displays **GET/** by default. You can specify an HTTP request up to 1024 characters. For more information, see [EditingHTTPRequestforHTTPHealthMonitor](#) below.
  - **Response Code Check:** Specify in which case the response code from the server is valid:
    - Select **Any response code is valid**, if any response code from the server is required.
    - Select **A valid response code**, select **equals** or **does not equal**, and then specify a value. The default value is 200.
  - **Response Content Check:** Specify an option for checking the server response content:
    - Select **Do not check the response content** to not perform any content check.
    - Select **Search for a string in the response content** to search for a string in the response content. Then do the following:
      1. In the **Search in** drop-down list, choose where to perform the search for a string: in **Both the header or body**, **Body**, or **Headers** of the HTTP request. The search is limited to the first five kilobytes of the response.
      2. In the **Regular Expression** field, specify a regular expression that will be used to search for a string in the response content.
      3. In the drop-down list **The content is valid if the regular expression is**, select either **found** or **not found**. If you select **found**, the content is valid if it corresponds to the regular expression you specify. If you select **not found**, the content is valid if it does not correspond to the regular expression you specify.
    - Select **Extract content from the response and compare it to a value** to extract a certain part of the content and compare it to a specific string or integer value. Then do the following:
      1. In the **Search in** drop-down list, choose where to perform the search for a string: in **Both the header or body**, **Body**, or **Headers** of the HTTP request. Note that the search is limited to the first five kilobytes.
      2. In the **Regular Expression** field, specify a regular expression for content extraction. The regular expression can contain subexpressions that you may specify in the next step. If you set **1st** as the value for **Check content that is extracted using the <...> subexpression**, the format of the **Regular Expression** must be `expression(sub_expression1)`. The number of subexpressions must be same as or can be more than the value you specify for **Check content that is extracted using the <...> subexpression**. For example, the format must be `expression(sub_expression1)(sub_expression2)` when the **Check content that is extracted using the <...> subexpression** is set to **2nd**, which means that there must be at least two subexpressions.
      3. Select **Check all extracted content** to or select **Check content that is extracted using the <...> subexpression** and choose a number of the subexpression from the drop-down list. You can choose from the first to the eighth subexpression previously defined in the **Regular Expression** field.
      4. In the field **The extracted content is valid if it is a**, select the expected data type of the extracted content, **string** or **integer**, and select a comparison operator. Then specify a value in the text field.
6. Click **Next** to add extensible attributes. For information, see [Using Extensible Attributes](#).
7. To schedule the change, click **Next** or **ScheduleforLater**. In the *ScheduleChange* panel, click **Now** to immediately execute this task. Or, click **Later** to schedule this task, and then specify a date, time, and time zone.
8. Save the configuration.

## Editing HTTP Request for HTTP Health Monitor

You can specify a multi-line message and include HTTP headers in the request by using the **HTTP Request** field in the HTTP health monitor properties. The header lines of an HTTP request have the simple name: value syntax. The request

headers are used to pass cookies, authentication, and provide information about the client to the server, etc. The HTTP 1.1 contains a request line and a single host header:

Example 1:

```
GET / HTTP/1.1
Host: www.yoursite.com
Connection: close
```

Example 2:

```
GET /index.html HTTP/1.1
Host: www.example.com
```

Note that the lines are terminated with two chars

`\r\n`. The whole request terminates with an empty line "`\r\n\r\n`" character sequence. NIOS adds `\r\n\r\n` string if it is absent. You can request `GET /` instead of "`GET/index.html`".

The host header differentiates between several HTTP servers that are running on a single IP address on the same port.

In an HTTP 1.1 request, the server keeps the connection alive by default after the response is sent. You can disable the connection by adding a `Connection: close` header line to the request.

An HTTP 1.0 request may consist of a single line followed by the automatically added `\r\n\r\n`:

```
GET /index.html HTTP/1.0
```

Or in the most simple form:

```
GET / HTTP/1.0
```

The server closes the connection after the response has been sent. You can use `Connection: Keep-Alive` header to alter this behavior. The `Content-Length` header is important to determine the end of the response for keep-alive connections.

Apart from HTTP 1.0/1.1, NIOS also supports a request format known as HTTP 0.9:

```
GET /index.html
```

or

```
GET /
```

Normally, the response header consists of a response line, such as `HTTP/1.1 200 OK` or `HTTP/1.0 400 Bad Request`, followed by a couple of header lines, and then by an empty line which signals the end of the response header. With HTTP 0.9, the response immediately starts with the content of the requested file, which means that there is no HTTP return code for an HTTP 0.9 request.

## Testing HTTP Health Monitors

After the HTTP health monitor is configured, you can test the configuration for a specific DTC server. Note that if you make changes to the HTTP health monitor settings, you must save the configuration so you can run the test.

To test the HTTP health monitor, do the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, and then click **Manage Health Monitors** in the Toolbar.
2. In the Health Monitors Manager, click the Action icon

- next to the HTTP health monitor name and select **Edit**.
3. Select the **Request/Response** tab.
  4. Click **Test HTTP Health Monitor**.
  5. In the field **Select a DTC Server or enter the IP address or domain name of an HTTP server**, do one of the following to specify the server to test:
    - Click **Select** to select an existing DTC server.
    - Enter the IP address or host name of an HTTP server. The IP address can be IPv4 or IPv6.
  6. In the field **Select a Grid member that is running DTC**, select a DTC server on which the test will be run. If there is only one DTC server with the DTC license, it is selected by default. If there are several DTC servers with the license, the Grid Master is selected by default. If there is no Grid Master with the DTC license and there are several member servers with the license, click **Select** and choose a server.
  7. Click **Test**.
  8. In the result of the test, the following information is returned:
    - a. Test status
    - b. Status message

**Note:** If you configure an HTTP or HTTPS monitor for a pool and try to test the monitor for HTTP version 1.1 using the **Test HTTP Health Monitor** option, you must add the host header for the health check to be successful.

## Configuring ICMP Health Monitors

An ICMP monitor sends an ICMP or ICMPv6 echo request to the IP address of the target server and expects an ICMP/ICMPv6 echo response. The ICMP monitor determines the health of a server by monitoring the response to an ICMP ping.

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, and then click **Manage Health Monitors** in the Toolbar.
2. In the Health Monitors Manager, click the arrow next to the Add icon and select **ICMP Health Monitor**.
3. In the *ICMP Health Monitor Wizard*, complete the following:
  - **Name:** Enter a name for the ICMP health monitor.
  - **Interval (seconds):** Enter the interval value in seconds. The health monitor runs only for the specified interval and it is measured from the end of the previous monitor cycle. The default value is five.
  - **Timeout (seconds):** Enter the timeout value in seconds. The monitor waits for the number of seconds that you specify after sending a response. If it does not receive a response within the number of seconds that you specify, then the appliance considers this check as failed. The monitor discards any response it receives after the timeout. The default value is 15.
  - **Retry Up Count:** Enter a retry up count integer value. Retry up count is a value that determines how many valid responses or good health checks in a row must be received by the Grid member from the DTC server for setting the DTC server health status to green. When you specify a value, the appliance computes the duration for which health check must be performed based on the following:  

$$\text{interval} * \text{retry up count}$$

For example, If the DTC server has had a red status for a long time because all the ICMP health checks have failed, and when the Grid receives the first good ICMP health check result, the health check counter is set as 1; however, the status of the DTC server is still retained as red. If you specify the **Retry Up Count** as 3 and the health check interval as 5, when three consecutive good health check results are received, the health check counter value becomes 3 and the DTC server health status now changes to green (3 good health check results in a row at an interval of 5 seconds each over a span of  $3 * 5 = 15$  seconds).
  - **Retry Down Count:** Enter a retry down count integer value. It is the opposite of the retry up count. Retry down count is a value that determines how many red status health checks (server is unavailable) in a row must be collected by the Grid member from the DTC server to switch the health status from green to red. For example, for a Grid member whose connection to the DTC server is not stable and has its Retry Down Count set as 5, if the health check results received by the Grid member is in the sequence: green, red, red, red, green, then the health status of the DTC server will be retained as green.

Note that red health status is set when the health monitor reaches the timeout value that is `[health check interval + timeout]` seconds without a valid response.

- **Comment:** Enter information about the ICMP health monitor.

4. Click **Next** to add extensible attributes. For information, see [Using Extensible Attributes](#).

5. To schedule the change, click **Next** or **Schedule for Later**. In the *Schedule Change* panel, select **Now** to immediately execute this task. Or select **Later** to schedule this task, and then specify a date, time, and time zone.

6. Save the configuration.

## Configuring PDP Health Monitors

A PDP (Packet Data Protocol) monitor sends a standard `GTP ECHO` request to the server. The GTP (GPRS Tunneling Protocol) echo message is used to ping the server. The connection is successful when the monitor receives an ECHO response from the server. If the server does not respond after a specified number of echo requests, the server is declared down by the monitor. You cannot modify the request or response.

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, and then click **Manage Health Monitors** in the Toolbar.
2. In the Health Monitors Manager, click the arrow next to the Add icon and select **PDP Health Monitor**.
3. In the *PDP Health Monitor Wizard*, complete the following:
  - **Name:** Enter a name for the PDP health monitor.
  - **Interval (seconds):** Enter the interval value in seconds. The health monitor runs only for the specified interval and it is measured from the end of the previous monitor cycle. The default value is five.
  - **Timeout (seconds):** Enter the timeout value in seconds. The monitor waits for the number of seconds that you specify after sending a response. If it does not receive a response within the number of seconds that you specify, then the appliance considers this check as failed. The monitor discards any response it receives after the timeout. The default value is 15.
  - **Retry Up Count:** Enter a retry up count integer value. Retry up count is a value that determines how many valid responses or good health checks in a row must be received by the Grid member from the DTC server for setting the DTC server health status to green. When you specify a value, the appliance computes the duration for which health check must be performed based on the following:

`interval*retry up count`

For example, If the DTC server has had a red status for a long time because all the PDP health checks have failed, and when the Grid receives the first good PDP health check result, the health check counter is set as 1; however, the status of the DTC server is still retained as red. If you specify the **Retry Up Count** as 3 and the health check interval as 5, when three consecutive good health check results are received, the health check counter value becomes 3 and the DTC server health status now changes to green (3 good health check results in a row at an interval of 5 seconds each over a span of  $3*5=15$  seconds).

- **RetryDownCount:** Enter a retry down count integer value. It is the opposite of the retry up count. Retry down count is a value that determines how many red status health checks (server is unavailable) in a row must be collected by the Grid member from the DTC server to switch the health status from green to red. For example, for a Grid member whose connection to the DTC server is not stable and has its Retry Down Count set as 5, if the health check results received by the Grid member is in the sequence: green, red, red, red, green, then the health status of the DTC server will be retained as green.

Note that red health status is set when the health monitor reaches the timeout value that is `[health check interval + timeout]` seconds without a valid response.

- **Port:** Specify a port for PDP connection. The appliance displays 2123 by default. You can specify a value between zero and 65535.
- **Comment:** Enter information about the PDP health monitor.

4. Click **Next** to add extensible attributes. For information about using attributes, see [Managing Extensible Attributes](#).

To schedule the change, click **Next** or **Schedule for Later**. In the *Schedule Change* panel, select **Now** to immediately execute this task. Or select **Later** to schedule this task, and then specify a date, time, and time zone.



5. Save the configuration.

## Configuring SIP Health Monitors

A SIP monitor sends a standard `SIP OPTIONS` request to the server. You cannot modify this request. The monitor accepts only direct responses from the server and does not open alternate connections. The SIP monitor determines the health of the SIP server such as SIP proxies and session border controllers, and SIP gateways by issuing `SIP OPTIONS` to the server and examining the response provided by the server. The service is considered available if the response received from the server matches the expected response.

The SIP monitor does not support SCTP transport. It does not receive SIP connections. Responses are normally received over the same connection as the request was sent. The server does not attempt to open a new connection to send the response when it encounters an error message.

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, and then click **Manage Health Monitors** in the Toolbar.
2. In the Health Monitors Manager, click the arrow next to the Add icon and select **SIP Health Monitor**.
3. In the *SIP Health Monitor Wizard*, complete the following:
  - **Name**: Enter a name for the SIP health monitor.
  - **Interval (seconds)**: Enter the interval value in seconds. The health monitor runs only for the specified interval and it is measured from the end of the previous monitor cycle. The default value is five.
  - **Timeout (seconds)**: Enter the timeout value in seconds. The monitor waits for the number of seconds that you specify after sending a response. If it does not receive a response within the number of seconds that you specify, then the appliance considers this check as failed. The monitor discards any response it receives after the timeout. The default value is 15.
  - **Retry Up Count**: Enter a retry up count integer value. Retry up count is a value that determines how many valid responses or good health checks in a row must be received by the Grid member from the DTC server for setting the DTC server health status to green. When you specify a value, the appliance computes the duration for which health check must be performed based on the following:  
`interval*retry up count`  
For example, If the DTC server has had a red status for a long time because all the SIP health checks have failed, and when the Grid receives the first good SIP health check result, the health check counter is set as 1; however, the status of the DTC server is still retained as red. If you specify the **Retry Up Count** as 3 and the health check interval as 5, when three consecutive good health check results are received, the health check counter value becomes 3 and the DTC server health status now changes to green (3 good health check results in a row at an interval of 5 seconds each over a span of 3\*5=15 seconds).
  - **Retry Down Count**: Enter a retry down count integer value. It is the opposite of the retry up count. Retry down count is a value that determines how many red status health checks (server is unavailable) in a row must be collected by the Grid member from the DTC server to switch the health status from green to red. For example, for a Grid member whose connection to the DTC server is not stable and has its Retry Down Count set as 5, if the health check results received by the Grid member is in the sequence: green, red, red, red, green, then the health status of the DTC server will be retained as green.  
Note that red health status is set when the health monitor reaches the timeout value that is `[health check interval + timeout]` seconds without a valid response.
  - **Comment**: Enter information about the SIP health monitor.  
When the DTC server is up, health checks are sent from the Infoblox member every number of seconds specified in the **Interval** field. When the DTC server is down, health checks are sent every number of seconds specified in the **Interval** field + the **Timeout** field.
4. Click **Next** and complete the following:
  - **Expected Return Code**: The response code expected from the server. Select a value from the drop-down list: **any**, **equals**, and **does not equals**. When you select **equals** or **does not equals**, the appliance displays 200 by default. You can specify a value between zero and 999.
  - **Port**: Specify a port for SIP connection. The appliance displays 5060 for **TCP** and **UDP** transport by default. When you select **SIPS** and **TLS** transport options, the appliance displays 5061 by default. You can specify a value between zero and 65535.



- **Transport:** Select a transport option from the drop-down list: **SIPS, TCP, TLS, and UDP**. If you select **SIPS** or **TLS**, specify any of the following related options that become available:
  - **Client Certificate:** Click **Certificate** to select a client certificate. Select a certificate from the dialog box. Click **Clear** to delete the certificate that you have uploaded. The monitor does not inspect or validate the server certificate, if any. For information about how to upload certificates, see *ManagingHealthMonitor Certificates* below.
  - **Ciphers:** Specify a list of SSL ciphers in an OpenSSL format. You can specify text up to 1024 character.
  - **Enable Certificate Validation:** It is highly recommended to select this for the DTC server certificate to be validated by NIOS.  
The following example commands list some available ciphers:

Example 1:

```
$ openssl ciphers 'HIGH:!DES'
DHE-RSA-AES256-SHA:DHE-DSS-AES256-SHA:DHE-RSA-CAMELLIA256-SHA:
DHE-DSS-CAMELLIA256-SHA:ADH-AES256-SHA:ADH-CAMELLIA256-SHA:AES256-SHA:
CAMELLIA256-SHA:PSK-AES256-CBC-SHA:EDH-RSA-DES-CBC3-SHA:
EDH-DSS-DES-CBC3-SHA:ADH-DES-CBC3-SHA:DES-CBC3-SHA:DES-CBC3-MD5:
PSK-3DES-EDE-CBC-SHA:KRB5-DES-CBC3-SHA:KRB5-DES-CBC3-MD5:DHE-RSA-AES128-SHA:
DHE-DSS-AES128-SHA:DHE-RSA-CAMELLIA128-SHA:DHE-DSS-CAMELLIA128-SHA:
ADH-AES128-SHA:ADH-CAMELLIA128-SHA:AES128-SHA:CAMELLIA128-SHA: PSK-AES128-CBC-
SHA
```

Example 2:

```
$ openssl ciphers 'DEFAULT:!EDH+aRSA'
DHE-DSS-AES256-SHA:DHE-DSS-CAMELLIA256-SHA:AES256-SHA:CAMELLIA256-SHA:
PSK-AES256-CBC-SHA:EDH-DSS-DES-CBC3-SHA:DES-CBC3-SHA:PSK-3DES-EDE-CBC-SHA:
KRB5-DES-CBC3-SHA:KRB5-DES-CBC3-MD5:DHE-DSS-AES128-SHA:DHE-DSS-SEED-SHA: DHE-
DSS-CAMELLIA128-SHA:AES128-SHA:SEED-SHA:CAMELLIA128-SHA:
PSK-AES128-CBC-SHA:RC4-SHA:RC4-MD5:PSK-RC4-SHA:KRB5-RC4-SHA:KRB5-RC4-MD5: EDH-
DSS-DES-CBC-SHA:DES-CBC-SHA:KRB5-DES-CBC-SHA:KRB5-DES-CBC-MD5:
EXP-EDH-DSS-DES-CBC-SHA:EXP-DES-CBC-SHA:EXP-RC2-CBC-MD5:EXP-KRB5-RC2-CBC-SHA:
EXP-KRB5-DES-CBC-SHA:EXP-KRB5-RC2-CBC-MD5:EXP-KRB5-DES-CBC-MD5:EXP-RC4-MD5: EXP-
KRB5-RC4-SHA:EXP-KRB5-RC4-MD5
```

Note

The DHE cipher list family ("Diffie-Hellman key agreement" plus "RSA authentication") could consume excessive CPU and is excluded from the defaults used by DNS Traffic Control health monitors. Although you can enable these ciphers by explicitly configuring them in the cipher list for HTTPS and SIP monitors, you should be aware that doing so will increase CPU usage. Since health monitoring in general does not require high security, Infoblox recommends that you enable these ciphers only for target servers that do not accept other types of ciphers.

5. Click **Next** to add extensible attributes. For information, see [Using Extensible Attributes](#).

6. To schedule the change, click **Next** or **Schedule for Later**. In the *Schedule Change* panel, select **Now** to immediately execute this task. Or select **Later** to schedule this task, and then specify a date, time, and time zone.

## 7. Save the configuration.

### Configuring SNMP Health Monitors

An SNMP health monitor sends an SNMPv1, SNMPv2c, or SNMPv3 request to the monitored server. The SNMP agent in the managed server provides the data in the form of variables, and each variable is associated with a unique OID (object identifier). An OID is a dotted-decimal number that defines the location of the object in the universal MIB tree. You can manually enter up to 15 OIDs to be monitored by the SNMP monitor. The server is considered available if the response received from the server matches the expected result for all OIDs. If the server does not respond after a specified number of requests, the server is declared down by the monitor.

To configure an SNMP OID health monitor, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, and then click **Manage Health Monitors** in the Toolbar.
2. In the Health Monitors Manager, click the arrow next to the Add icon and select **SNMP Health Monitor**.
3. In the *SNMP Health Monitor Wizard*, complete the following:
  - **Name:** Enter a name for the SNMP monitor.
  - **Interval (seconds):** Enter the interval value in seconds. The health monitor runs only for the specified interval and it is measured from the end of the previous monitor cycle. The default value is five.
  - **Timeout (seconds):** Enter the timeout value in seconds. The monitor waits for the number of seconds that you specify after sending a response. If it does not receive a response within the number of seconds that you specify, then the appliance considers this check failed. The monitor discards any responses it receives after the timeout. The default value is 15.
  - **Retry Up Count:** Enter a retry up count integer value. Retry up count is a value that determines how many valid responses or good health checks in a row must be received by the Grid member from the DTC server for setting the DTC server health status to green. When you specify a value, the appliance computes the duration for which health check must be performed based on the following:  
$$\text{interval} \times \text{retry up count}$$

For example, If the DTC server has had a red status for a long time because all the SNMP health checks have failed, and when the Grid receives the first good SNMP health check result, the health check counter is set as 1; however, the status of the DTC server is still retained as red. If you specify the **Retry Up Count** as 3 and the health check interval as 5, when three consecutive good health check results are received, the health check counter value becomes 3 and the DTC server health status now changes to green (3 good health check results in a row at an interval of 5 seconds each over a span of  $3 \times 5 = 15$  seconds).
  - **Retry Down Count:** Enter a retry down count integer value. It is the opposite of the retry up count. Retry down count is a value that determines how many red status health checks (server is unavailable) in a row must be collected by the Grid member from the DTC server to switch the health status from green to red. For example, for a Grid member whose connection to the DTC server is not stable and has its Retry Down Count set as 5, if the health check results received by the Grid member is in the sequence: green, red, red, red, green, then the health status of the DTC server will be retained as green.  
  
Note that red health status is set when the health monitor reaches the timeout value that is  $[\text{health check interval} + \text{timeout}]$  seconds without a valid response.
  - **Port:** Specify a port for the SNMP connection. The appliance displays 161 by default. You can specify a value between zero and 65535.
  - **Comment:** Enter information about the SNMP health monitor.
4. Click **Next** and complete the following:
  - **Version:** Select the SNMP version, **v1**, **v2c**, or **v3**. Note that the available options for versions v1 and v2c differ from those for v3 version.
  - (SNMPv1 and SNMPv2c only) **Community:** Enter the text string that the SNMP monitor must send along with the queries to the server for authentication. The community string is similar to a password and the server accepts queries only from the SNMP monitor that provide the correct community string. Note that this community string must match exactly what you enter in the management system. The default value is **public**.
  - (SNMPv3 only) **SNMPv3User:** Click **Select** or **Create** to specify an SNMPv3 user. For information about SNMPv3 users, see [Configuring SNMP](#).

If you are modifying an already existing SNMPv3 health monitor in the SNMP Health Monitor editor, two additional optional fields become available:

- **Context:** enter an arbitrary string.
- **EngineID:** enter an arbitrary string that can contain from 10 to 64 hexadecimal digits (5 to 32 octet numbers).

Click the Add icon above the **Health Monitor SNMP OIDs** table to add an SNMP OID entry. Complete the following:

- **OID:** Specify the object identifier. An OID is a unique dotted-decimal number that identifies the location of the object in the MIB tree. For more information about OIDs, see [SNMP MIB Hierarchy](#).
- **Type:** Select either **String** or **Integer** from the drop-down list.

**Note:** If you use this SNMP monitor with the Ratio: Dynamic load balancing method, note that only integer OID type is supported for this method.

- **Operator:** Select one of these operators from the drop-down list: **Any**, **Equals**, **Larger or equals**, **Range**, and **Smaller or equals**.
- **Value:** If the operator is **Equals**, **Larger or equals**, or **Smaller or equals**, enter a value. If the operator is **Range**, enter the minimum and maximum values in the **Min value** and **Max value** fields respectively.
- **Comment:** Enter information about the SNMP OID entry.
- Click **Add** to add the SNMP OID to the table.

5. Click **Next** to add extensible attributes. For information, see [Using Extensible Attributes](#).

6. To schedule the change, click **Next** or **Schedule for Later**. In the *Schedule Change* panel, select **Now** to immediately execute this task. Or select **Later** to schedule this task, and then specify a date, time, and time zone.

7. Save the configuration.

## Configuring TCP Health Monitors

A TCP monitor opens a TCP connection to communicate between the appliance and server. The connection is successful only when the handshake is complete. A successfully opened connection will be immediately closed or reset.

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, and then click **Manage Health Monitors** in the Toolbar.
2. In the Health Monitors Manager, click the arrow next to the Add icon and select **TCP Health Monitor**.
3. In the *TCP Health Monitor Wizard*, complete the following:
  - **Name:** Enter a name for the TCP health monitor.
  - **Interval (seconds):** Enter the interval value in seconds. The health monitor runs only for the specified interval and it is measured from the end of the previous monitor cycle. The default value is five.
  - **Timeout (seconds):** Enter the timeout value in seconds. The monitor waits for the number of seconds that you specify after sending a response. If it does not receive a response within the number of seconds that you specify, then the appliance considers this check as failed. The monitor discards any response it receives after the timeout. The default value is 15.
  - **Retry Up Count:** Enter a retry up count integer value. Retry up count is a value that determines how many valid responses or good health checks in a row must be received by the Grid member from the DTC server for setting the DTC server health status to green. When you specify a value, the appliance computes the duration for which health check must be performed based on the following:

`interval*retry up count`

For example, If the DTC server has had a red status for a long time because all the TCP health checks have failed, and when the Grid receives the first good TCP health check result, the health check counter is set as 1; however, the status of the DTC server is still retained as red. If you specify the **Retry Up Count** as 3 and the health check interval as 5, when three consecutive good health check results are received, the health check counter value becomes 3 and the DTC server health status now changes to green (3 good health check results in a row at an interval of 5 seconds each over a span of  $3*5=15$  seconds).

- **Retry Down Count:** Enter a retry down count integer value. It is the opposite of the retry up count. Retry down count is a value that determines how many red status health checks (server is unavailable) in a row must be collected by the Grid member from the DTC server to switch the health status from green to red. For example, for a Grid member whose connection to the DTC server is not stable and has its Retry Down

Count set as 5, if the health check results received by the Grid member is in the sequence: green, red, red, red, green, then the health status of the DTC server will be retained as green.

Note that red health status is set when the health monitor reaches the timeout value that is `[health check interval + timeout]` seconds without a valid response.

- **Port:** Specify a port for TCP connection. You can specify a value between zero and 65535.
  - **Comment:** Enter information about the TCP health monitor.
4. Click **Next** to add extensible attributes. For information, see [Using Extensible Attributes](#).
  5. To schedule the change, click **Next** or **Schedule for Later**. In the *Schedule Change* panel, select **Now** to immediately execute this task. Or select **Later** to schedule this task, and then specify a date, time, and time zone.
  6. Save the configuration.

## Managing Health Monitor Certificates

You can upload multiple certificates to the appliance and associate them with HTTP and SIP health monitors. The appliance supports certificates that are in PEM or PKCS#12 format only. A PEM file can contain more than one certificate. Note that the uploaded certificate must include both the client certificate and the private key. You can add, delete or view certificates.

To upload a health monitor certificate:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, and then click **Manage Health Monitors** in the Toolbar.
2. In the Health Monitors Manager, click the Manage Health Monitor Certificates icon.
3. In the **Health Monitor Certificates** window, click the Add icon.
4. In the **Upload** window, click **Select** and navigate to the certificate you want to upload.
5. Select the file and click **Upload**.

Grid Manager displays the following information in the **Health Monitor Certificates** window:

- **Issuer:** The name of the trusted CA that issued the certificate.
- **Valid From:** The date from which the certificate becomes valid.
- **Valid To:** The date until which the certificate is valid.
- **Subject:** The name of the certificate.

To upload a certificate from a web browser:

1. Export a PEM file from a web browser.
2. Generate a private key using the following command:

```
- openssl req -x509 -nodes -days 365 -newkey rsa:1024 -keyout mycert.pem -out mycert.pem
```
3. Copy the private key section out of mycert.pem and add it to the PEM file from the web browser.
4. Or add both the PEM and the private key together as PKCS#12.

You can also do the following in the **Health Monitor Certificates** window:

- Click the checkbox next to the issuer and click the Delete icon to delete it.
- Print the data or export it in .csv format.

## Modifying Health Monitors

To modify a health monitor:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, and then click **Manage Health Monitors** in the Toolbar.
2. In the Health Monitors Manager, click the Action icon next to the health monitor name, and select **Edit** from the menu.
3. The *Health Monitor* editor contains the following tabs from which you can modify data:
  - **General:** This tab displays the health monitor **Name**, **Comment**, **Interval**, **Timeout**, **Retry Up Count**, and **Retry Down Count** fields. You can edit the values.
  - (HTTP, SIP, and SNMP health monitors) **Protocol:** This tab displays the protocol data that you can modify.

- (HTTP/S health monitors only) **Request/Response**: This tab displays HTTP request and response check options for HTTP/S health monitors.
- **Extensible Attributes**: Add and delete extensible attributes that are associated with the pool. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).

For information about modifying the details of specific monitors, see the corresponding sections above.

4. To schedule this task, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, select **Later** and enter a date, time, and time zone. The Schedule icon is green when there is a pending scheduled task. For information, see [Scheduling Tasks](#).

5. Save the configuration.

## Viewing Health Monitors

You can view health monitors that you have created. You can add new health monitors, delete existing monitors, modify health monitors, or associate extensible attributes to them. You can also upload and manage health monitor certificates. To view health monitors:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, and then click **Manage Health Monitors** in the Toolbar.
2. The Health Monitors Manager displays the following information:
  - **Name**: The name of the health monitor.
  - **Type**: The type of health monitor.
  - **Comment**: Displays information about the health monitor.
  - **Interval**: The interval value in seconds.
  - **Timeout**: The timeout value in seconds.
  - **Retry Up Count**: The retry up count specified for the respective health monitor.
  - **Retry Down Count**: The retry down count specified for the respective health monitor.
  - **Port**: The port number specified for the respective health monitor. Note that this is not valid for an ICMP monitor.
  - **Site**: Value that was entered for the respective health monitor.

You can do the following in the Health Monitors Manager:

- Define new health monitors. For more information, see the following sections:
  - [ConfiguringHTTPHealthMonitors](#)
  - [ConfiguringICMPHealthMonitors](#)
  - [ConfiguringPDPHealthMonitors](#)
  - [ConfiguringSIPHealthMonitors](#)
  - [ConfiguringSNMPHealthMonitors](#)
  - [ConfiguringTCPHealthMonitors](#)
- Edit existing health monitors.
- Manage health monitor certificates.
- Delete existing health monitors.
- Click the Export icon to export the list of monitors to a .csv file.
- Click the Print icon to print the list of monitors.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the monitor from the possible matches. Create a quick filter to save frequently used filter criteria. For information about using quick filters, see [Finding and Restoring Data](#).

## Deleting Health Monitors

To delete a health monitor:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab, and then click **Manage Health Monitors** in the Toolbar.
2. In the Health Monitors Manager, click the Action icon next to the health monitor name and select **Delete** from the menu.

3. In the *Delete Confirmation* dialog box, click **Yes** to delete or **No** to cancel.

Click **Schedule Deletion** and in the Schedule Change panel, enter a date, time, and time zone to schedule deletion at a later date and time. For more information about schedule deletion, see [Scheduling New IPAM/DHCP Objects and Associated Port Configurations](#).

## Configuring DTC Monitors for Health Check

You can choose DNS Traffic Control health monitors whose DNS Traffic Control health checks are considered when calculating the health status of a member in a pool. You can create a consolidated health monitor in a DNS Traffic Control pool with its own set of members for each monitor. The status of a health check can then be determined according to the consolidated health monitor you created with selected members and availability condition.

To configure DNS Traffic Control monitors whose health check status must be shared from selected members, perform the following steps:

1. From the **Data Management** tab, select the **DNS** tab -> **Traffic Control** tab.
2. Select the pool for which you want to configure monitors and click the **Edit** icon.
3. Click the **Health Monitors** tab -> **Advanced** tab.
4. Click the **Add** icon.
5. From the **Monitor** drop-down list, choose the monitor whose health checks you want to be considered.
6. From the **Availability Requirement** drop-down list, you can either select all members of the pool to pass the health check for the selected monitor or for any of the pool member to pass the health check for the selected monitor. A passed health check is denoted by the green status.
7. From the **Share state from** box, select the members that you want the health check to be calculated. Select the members and move them to the **Selected** box.
8. Select the **Full Health Communication** checkbox to perform health checks on each and every member of the Grid which are part of LBDN and to send the status to all the Grid members which are in non selected list and the status is shared between grid members which are in selected list. This checkbox is disabled by default. By default, NIOS performs health checks only on Grid members from the selected list.
9. Click **Add**.  
The monitor along with its associated members and full health communication details are displayed in the **Consolidated Health Monitor Settings** area.
10. Click **Save & Close**.

### Note

- To add DNS Traffic Control pool with consolidated monitors using CSV import, you must first import the DNS Traffic Control pool objects without consolidated monitor data (if the CSV file contains consolidated monitors, you must manually remove the data before importing the DNS Traffic Control pool objects). Once the DNS Traffic Control pool objects are in the system, you can run the CSV override to add consolidated monitors using a CSV file that contains all the data, including the DNS Traffic Control pool objects and consolidated monitors.
- An error is displayed if there are no members added to consolidated health monitor, as the health status of DNS Traffic Control pool objects is not set properly.

You can now view a snapshot of any member in the pool and the monitors associated with it, along with their health status. To view the snapshot, complete the following:

1. Click the server for which you want to see the consolidated monitors.
2. In the **Server Visualization** area, select the member for which you want to see the associated monitors.

A visualization chart is displayed that represents the pool hierarchy. You can hover over the server icons to view the health status of the monitors associated with the server. For more information about health monitors, see [Using DNS Traffic Control Health Monitors](#).



## Configuring IP Routing Options

You can configure multiple IP addresses and enable anycast addressing on the loopback interface of the NIOS appliance, allowing the appliance to function in different network deployments.

Configuring non-anycast IP addresses on the loopback interface assists in server migration and network address change. Configuring anycast addresses on the appliance allows you to add redundancy and improve reliability for DNS services. You can use OSPF (Open Shortest Path First), BGP (Border Gateway Protocol), or both, as the routing protocol for anycast advertising.

- [Configuring IP Addresses on the Loopback Interface](#)
- [About Anycast Addressing for DNS](#)
- [Configuring Anycast Addresses](#)
- [IP Routing Options](#)
- [About BFD \(Bidirectional Forwarding Detection\)](#)

## Configuring IP Addresses on the Loopback Interface

The loopback interface is a virtual network interface on the appliance. You can do the following on the loopback interface:

- Configure IP addresses to consolidate DNS servers for migration purposes. For information, see [Configuring IP Addresses](#).
- Add anycast addresses to improve the reliability and performance of DNS services in multiple locations. For information, see [About Anycast Addressing for DNS](#).
- Separate DNS traffic by assigning an IP address as the source port for DNS queries. For information, see [Specifying Source Ports](#).

When you use the loopback interface for anycast addressing, the upstream and neighboring routers can continue to advertise anycast addresses without being affected by hardware malfunctions.

To configure non-anycast addresses on the loopback interface, complete the following:

1. Add IP addresses to the loopback interface. For information, see [Configuring IP Addresses](#).
2. Enable DNS services on the loopback addresses. For information, see [Specifying Port Settings for DNS](#) and its subtopic, [Specifying Source Ports](#).

To configure DNS anycast addresses and their advertising protocols, complete the following:

1. Add anycast addresses to the loopback interface. For information, see [Configuring Anycast Addresses](#).
2. Configure anycast addressing protocols. For information about Configuring OSPF on the NIOS Appliance and Configuring BGP in the NIOS Appliance, see [IP Routing Options](#). This is the primary application for routing protocols in the NIOS appliance.
3. Enable the DNS anycast addresses. For information, see [Specifying Port Settings for DNS](#) and its subtopic, [Specifying Source Ports](#).

To separate DNS queries from DNS transfers and notify messages, complete the following:

1. Add an IP address of the source port for DNS queries. For information, see [Configuring IP Addresses](#).
2. Select the source IP for DNS queries. For information, see [Specifying Source Ports](#).

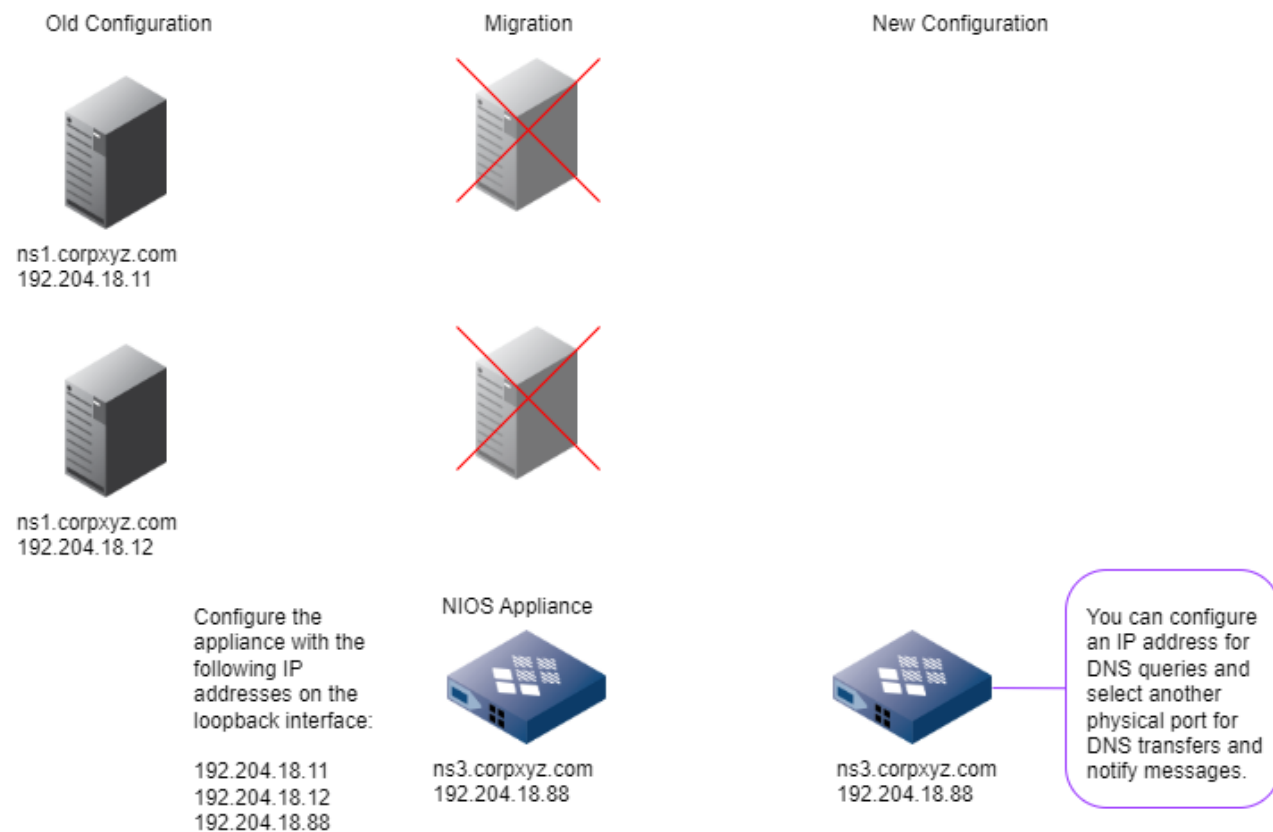
## Configuring IP Addresses

You can configure IP addresses on the loopback interface to minimize service downtime during a server migration. As illustrated in the figure DNS Server Migration Using the Loopback Interface below, you have two existing DNS servers ([ns1.corpxyz.com](#) 192.204.18.11 and [ns2.corpxyz.com](#) 192.204.18.12) and you want to replace these servers with a new one ([ns3.corpxyz.com](#) 192.204.18.88). The migration takes a few weeks and you want DNS services to be available on all three addresses during the migration. You can add all three IP addresses to the loopback interface of a



NIOS appliance, and then configure the appliance to provide DNS services on all addresses. After the server migration, you can shut down the old servers and use the new one for services.

### DNS Server Migration Using the Loopback Interface



You can also add an IP address that is used solely for DNS queries, to separate the DNS traffic. You first add an IP address you want to use for DNS queries on the loopback interface. You then configure the appliance to listen for DNS queries solely on this address. For information, see [Specifying Source Ports](#). When you configure non-anycast addresses on the loopback interface, ensure that you establish a static route between the appliance and the router so queries to these addresses are routed correctly. For information, see [Advertising Loopback Addresses to the Network](#).



#### Note

You can configure multiple interfaces on the Infoblox-4030-10GE appliance only. To configure LAN1, LAN2 and MGMT interfaces to the same IPv4 or IPv6 subnet, provide the same netmask for IPv4, or a CIDR prefix for IPv6, as the LAN1 interface. Alternatively, you can use a /32 netmask (255.255.255.255) for IPv4, or /128 CIDR prefix for IPv6 with the same subnet as LAN1 interface to configure multiple interfaces. An Infoblox-4030-10GE can replace three DNS cache servers that are active on the same network. When you configure multiple interfaces on the same subnet, the outgoing traffic from NIOS host which is received through LAN2 and MGMT is directed to the LAN1 router for all interfaces on the LAN1 subnet, irrespective of the destination IP. However, if the LAN1 interface fails, the outgoing traffic will not be re-directed to any other interface and access to LAN2 and MGMT also fails.

To configure an IP address on the loopback interface:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox -> Edit icon.
2. In the *Grid Member Properties Editor*, select the **Network** tab -> **Basic** tab.  
You can add an IPv4 or IPv6 address on the loopback. You define each type in their own table.

- Click the Add icon in the Additional Ports and Addresses table and select **Additional Address (loopback) (IPv4)** or **Additional Address (loopback) (IPv6)** from the drop-down list.

You cannot configure **Additional Address (loopback) (IPv4)** interface for an IPv6 Grid member and **Additional Address (loopback) (IPv6)** interface for an IPv4 Grid member. You can only enter the IP address you want to add to the loopback interface. You cannot configure the subnet mask, prefix length, gateway, or port settings.

The appliance adds a row to the table. Complete the following:

- Interface:** Displays **Additional Address (loopback)**. You cannot modify this.
  - Address:** Enter the IP address you want to add to the loopback interface. An IPv6 address is a 128-bit number in colon hexadecimal notation. It consists of eight 16-bit groups of hexadecimal digits separated by colons (example: 2001:db8:0000:0123:4567:89ab:0000:cdef or 2001:db8::123:4567:89ab:0:cdef).
  - Subnet Mask (IPv4) or Prefix Length (IPv6):** You cannot change the netmask of the loopback interface. It is set to 255.255.255.255, or /32. For an IPv6 address, the mask is set to 128 and cannot be modified. You cannot configure the gateway address and port settings.
- Save the configuration and click **Restart** if it appears at the top of the screen. To add multiple IP addresses on the loopback interface, repeat the steps for each IP address.



#### Note

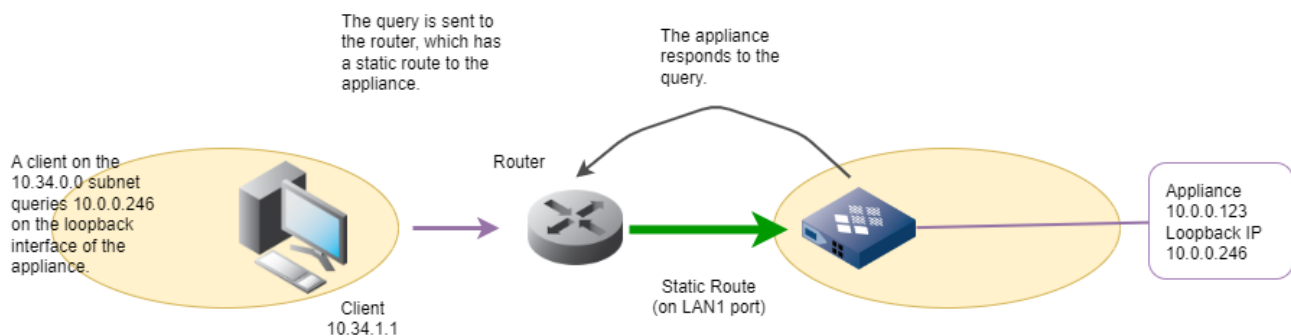
If you are configuring the loopback interface on a Grid Master, the Grid is temporarily disrupted upon saving the configuration and restarting services on the appliance. The Grid reconnects automatically and the appliance regains the role as Grid Master after a short delay.

## Advertising Loopback Addresses to the Network

Advertising IP addresses on the loopback interface relies on the upstream router to populate routes to the loopback interface. As illustrated in the figure Static Route for Loopback IP Addresses below, when a client on a different subnet queries an IP address on the loopback interface, it sends the request to the router. If the IP address on the loopback interface is not advertised to the router, the request cannot reach the appliance. Therefore, when you configure non-unicast addresses on the loopback interface, or if OSPF or BGP is not configured within your network, you must configure the upstream router to reach the NIOS appliance through a static route on the LAN1 interface.

Note that when an appliance is configured for both authoritative and recursive queries, you should connect your internet interface through the LAN1 port to allow for maximum flexibility while using auxiliary LAN2 and MGMT ports. Consult with your network administrator for information about configuring static routes from the router to the additional IP addresses on the loopback interface.

### Static Route for Loopback IP Addresses





#### Note

To enable the loopback address functionality on vNIOS for Azure instances, you must configure the following settings in the Microsoft Azure or the Microsoft Azure Stack Hub portal according to the environment in which you have deployed the vNIOS instance:

- Add an inbound port rule for the vNIOS VM on the *virtual machine* > **Settings** > **Networking** tab to allow inbound packets to the loopback address interface on port 53 (for DNS).
- Add a route to your subnets. To add, create a route table (if the subnet does not have one already) and add the static route to the loopback address in the route table. Ensure to associate the route table with the subnet if you created a new route table.”

When you configure DNS anycast addresses on the loopback interface, you can select OSPF, BGP, or both, to advertise the addresses to upstream and neighboring routers, without establishing a static route. For information, see [About Anycast Addressing for DNS](#).

## About Anycast Addressing for DNS

Four types of communications are utilized within an IP network:

- **Unicast** describes a one-to-one network communication between a single sender and a single recipient. The routing protocol determines the path through the network from the sender to the recipient based on the specific protocol or routing scheme. Unicast also describes the address type assigned to the recipient.
- **Multicast** describes a one-to-many network communication between a single sender and a specific group of recipients. All members within the group are intended recipients and each member receives a copy of the data from the sender. Multicast also describes the address type assigned to a group of recipients, used by the routing protocol to determine the path to the group.
- **Broadcast** is similar to multicast, the exception being that data is sent to every possible destination regardless of the groups or subnetwork. There is no specific group of recipients.
- **Anycast** describes a one-to-nearest communication between a single sender and the nearest recipient within a group. The routing protocol chooses one recipient within a target group based on the routing algorithm for the specific protocol, and sends data to that recipient only.

The NIOS appliance provides the following support for DNS anycast addressing:

- You can configure up to 20 anycast IP addresses on the loopback interface of each Grid member.
- Anycast IP addresses can be in IPv4 or IPv6 format. For all anycast IP addresses, the subnet mask value is always set to /32 for an IPv4 anycast IP or 128 for a 128-bit IPv6 address. These values are separate and distinct from the IP configuration on the NIOS appliance LAN port.
- The appliance advertises routing information of the anycast addresses through OSPF or BGP, or (seldom) both, depending on the deployment. Routers use the configured routing protocols to determine the best path to the nearest server. The appliance advertises the route information to the upstream or neighboring router, a router that forwards data on the network link and determines the forwarding path to destinations. For information, see [IP Routing Options](#).
- The appliance advertises and withdraws route information based on reachability information to DNS servers sent by the IP route advertisements.
- When you configure DNS anycast addressing on an appliance and use it as an NTP server, the appliance can answer NTP requests through the anycast IP address. For information about how to configure an appliance as an NTP server, see [Configuring a NIOS Appliance as an NTP Server](#).

Anycast addressing for DNS provides the following benefits:

- **Improved Reliability:** Anycast provides improved reliability because DNS queries are sent to an anycast IP address that is defined on multiple DNS servers in the NIOS Grid. If the nearest server somehow goes offline, then the router forwards the request to the next nearest DNS server advertising the target anycast IP address (see [Anycast Addressing for DNS Using OSPF](#) for an example).
- **Load Distribution:** Anycast distributes the load across multiple DNS servers based on network topology.

- **Improved Performance:** The NIOS appliance uses OSPF or BGP, depending on your configuration, to advertise anycast routing information to the upstream and neighboring routers. The routers determine the best route to the nearest DNS server. Anycast enables the queries to reach the nearest server more quickly, providing faster responses to DNS queries.



#### Note

For more information about anycast addressing, refer to *RFC 1546 "Host Anycasting Service"*.

## Configuring Anycast Addresses

Anycast addressing is supported on loopback interfaces on the NIOS appliance. IP configuration must be defined on the LAN1 interface before configuring DNS anycast addresses. Before creating IPv6 anycast IPs on the loopback interface, IPv6 must be enabled and configured on the LAN1 interface for the NIOS appliance, including the correct IPv6 gateway IP address.



#### Note

When you add an anycast address, you need to start the service for the advertising to take place. However, when you remove an anycast address, no service restart is required to stop the service. Anycast advertising stops immediately.

To enable and configure anycast addressing:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox -> Edit icon.
2. In the *Grid Member Properties* editor, Click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Anycast** tab.
4. Click the Add icon and choose **IPv4 Address** or **IPv6 Address**.
5. In the **Anycast Interfaces** list, enter the values or select the options for the new entry:
  - **Anycast Interface:** Anycast addressing is supported on the **loopback** only. This value is filled in automatically.
  - **Address:** Enter the IP address you want to assign as the anycast address to the loopback interface. Specify an IPv4 Address or an IPv6 Address based on the chosen type of address.
  - **Subnet Mask:** You cannot change the subnet mask of a loopback interface. The netmask is automatically set to 255.255.255.255, or /32; or 128 for IPv6.
  - **OSPF:** Select if you want the appliance to use OSPF to advertise the anycast address, and if necessary configure the OSPF settings. For information, see [Configuring OSPF on the NIOS Appliance](#). IPv4 and IPv6 options are configurable for this protocol. This is supported only for IPv4 and dual mode appliances, but not for IPv6 appliances.
  - **BGP:** Select this if you want the appliance to use BGP to advertise the anycast address, and then configure the BGP settings. This is supported only for IPv4 and dual mode appliances, but not for IPv6 appliances.  
You must configure at least one routing method for DNS anycast. You can configure OSPF, BGP, or both (in most cases only one protocol will be used). The appliance cannot save the anycast address if you do not complete at least one routing configuration. Anycast cannot be used without dynamic routing.
  - **Comments:** Enter a text string to help identify this interface and IP address.
6. *If using OSPF for the current appliance:* Under **OSPF Area Configuration**, click the Add icon. A new configuration block appears in the properties editor.
  - Enter the values for the OSPF configuration as described in the section [Configuring OSPF on the NIOS Appliance](#).
  - Click the **Add** down arrow icon in the **OSPF Area Configuration** section. The new OSPF configuration is saved into a table.
7. *If using BGP for the current appliance:* In the properties editor, scroll down to the configuration block for **BGP Configuration**. For information, see [Configuring BGP in the NIOS Appliance](#).

- In the **ASN** field, enter the Autonomous System ID number in which the NIOS appliance resides.
  - If necessary, modify the **BGP Timer Keep Alive** and **Hold Down** values. In most circumstances these values should be left at their defaults. Check your network's defined policies for the desired values if necessary.
  - Click the Add icon.
  - Enter the **Neighbor Router** IP address. This can be an IPv4 address or an IPv6 address.
  - Enter the **Remote ASN** (Autonomous System ID number) for the adjacent router.
8. Save the configuration. The system will warn that you must restart the appliance services in order to use the new configuration.
  9. Log back in to the appliance.
  10. From the **Data Management** tab, select the **DNS** tab -> **Members/Servers** tab -> *Grid\_member* checkbox -> Edit icon.
  11. Select **Toggle Advanced Mode** (if necessary), click **General** and the **Advanced** tab.
  12. Under **Listen on these additional IP addresses**, click the Add button. The list of one or more previously created IPv4 and IPv6 addresses for the loopback interfaces (created in Step 4) appear in this table. (If the Add button is not active here, this indicates that you have not configured any loopback interfaces with their IP addresses.) Should you need to configure other DNS properties on this page, see the topics in [Configuring the Grid to provide DNS Services](#).  
Note that if you remove an IP address from the list of Listen on these additional IP addresses, anycast advertising will stop immediately. No service restart is required to stop anycast from listening on this IP address.
  13. Click **Save and Close**.

Configured anycast interfaces are now enabled to carry DNS traffic. For further information, see [Specifying Source Ports](#).



#### Note

You can select different options for the restart sequence for anycast service with DNS, when the DNS restart is invoked. You can manage this sequence with the help of CLI commands.

For more information on the CLI commands, see

- [set restart\\_anycast\\_with\\_dns\\_restart](#)
- [show restart\\_anycast\\_with\\_dns\\_restart](#)

## Best Practices for Configuring Anycast Addresses

Infoblox highly recommends that you do the following before you configure an anycast address:

- Enable the anycast feature in the NIOS application.
- Install a valid DNS license, enable DNS and ensure that the DNS service is active.
- When you configure OSPF or OSPFv6, ensure that the OSPF monitor runs every four seconds.
- You must configure an IP address on the loopback interface.

## IP Routing Options

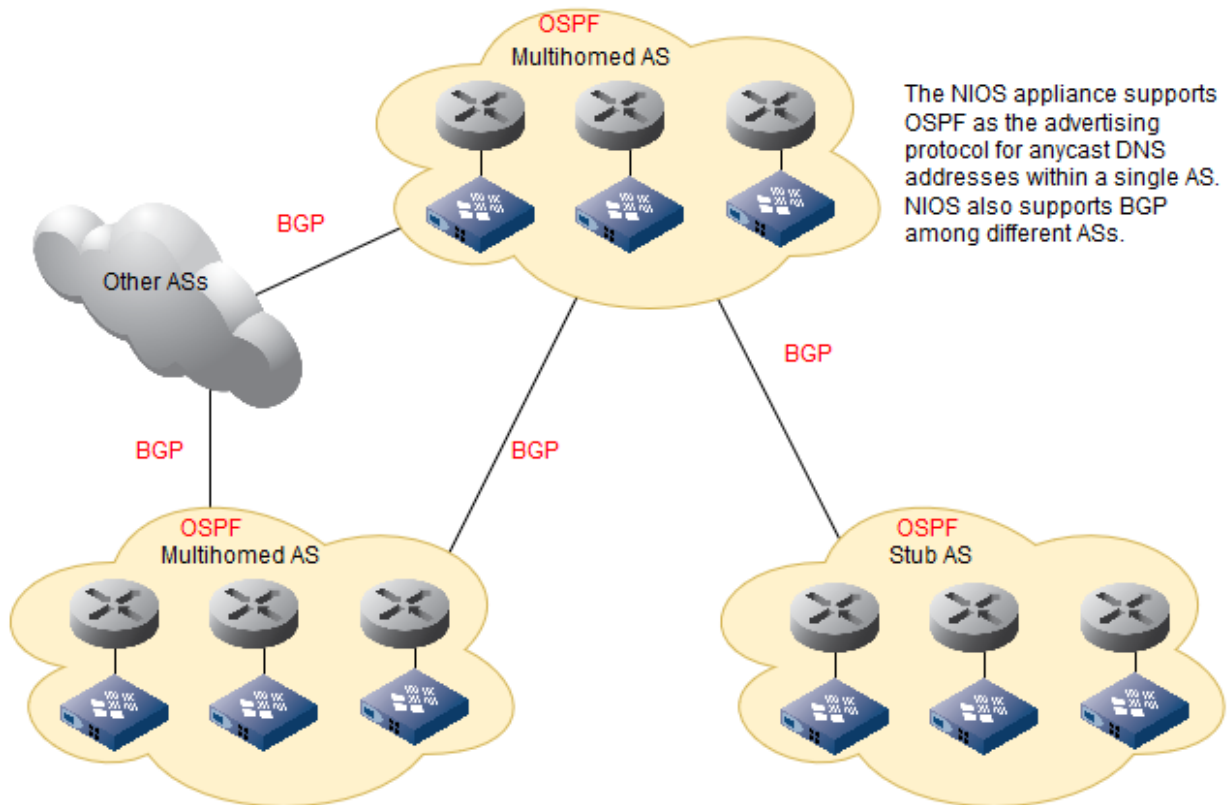
IP routing is a set of protocols that determine the path IP packets follow in order to travel across multiple networks from the source to the destination. When information travels through a series of routers and across multiple networks, IP routing protocols enable the routers to build up a forwarding table that correlates the final destination with the next upstream routers.

For routing purposes, the internet is divided into ASs (Autonomous Systems). Data is routed within an AS using an IGP (Interior Gateway Protocol) and routed between different ASs using an EGP (Exterior Gateway Protocol). NIOS appliances support OSPFv2 (for IPv4) and OSPFv3 (for IPv6) for a routing IGP, and BGP4 to advertise DNS anycast addresses in the larger internetwork.

As noted in the section [Configuring Anycast Addresses](#), you configure OSPF or BGP4 to advertise anycast addresses, which configured on the loopback interface of NIOS appliances. Use of either protocol depends on the

network topology, based on whether the advertisements will propagate only within a single AS or between more than one AS. The following figure shows a simplified example.

### OSPF and BGP Routing Example



Within each AS, OSPF is the protocol used to forward anycast advertisements. Between ASs, BGP is the protocol selected to advertise anycast addresses. Using this technique, DNS servers in diverse locations can operate together to ensure continuous service.

### About OSPF

OSPF is a link-state protocol based on the Dijkstra algorithm used to calculate the shortest path to a destination address within an internetwork. This protocol uses a link-state database created using routing information advertised from neighbors and peers, each with costs based on the state of that link to the destination. OSPF network topologies consist of administrative domains called OSPF areas. An area is a logical collection of OSPF routers, servers and other network devices that have the same area identifier. A router within an area keeps an OSPF database for its OSPF area only, reducing the size of the database that is maintained.

### Anycast and OSPF

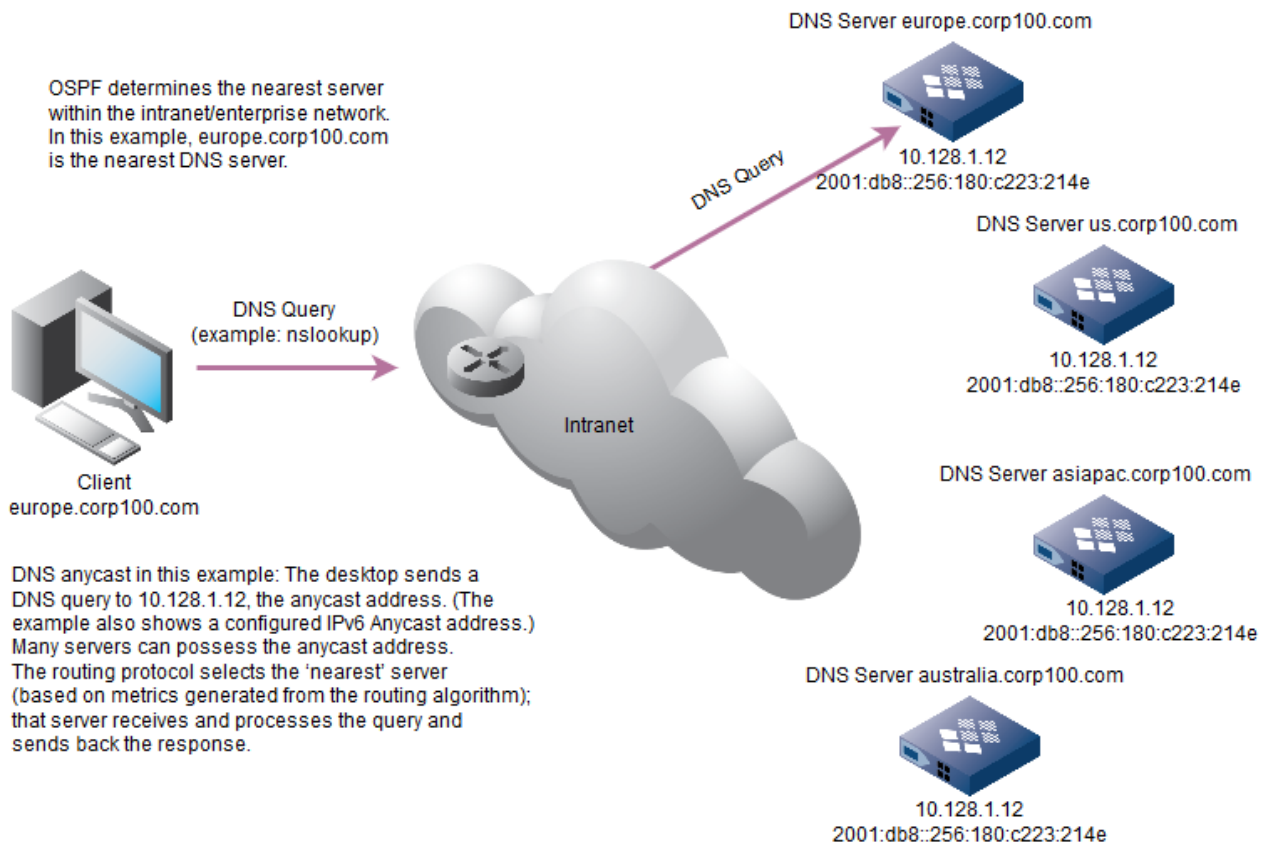
NIOS appliances can use the OSPF routing protocol to advertise routes for DNS anycast addresses to an upstream router within the autonomous system. The upstream router uses the OSPF advertisement to determine the nearest DNS server from a group of servers within the internetwork. In practice, the NIOS appliance relies upon OSPF to determine the best route for DNS queries to take to the nearest DNS server. The upstream router then forwards the query to the chosen DNS server.

As illustrated in the figure Anycast Addressing for DNS Using OSPF below, to enable anycast for DNS queries, you configure two or more DNS servers within the AS routing domain with the same anycast address on their loopback interfaces. When you select OSPF as the routing protocol, the upstream router determines the nearest server within the group of servers configured with that anycast address. (The "nearest" DNS server may not necessarily be the



geographically closest DNS server; it is the DNS server with the lowest cost associated with its reachability from the current node. This is calculated through the OSPF routing algorithm, a discussion of which is far beyond the scope of this manual.) The nearest DNS server configured with the correct anycast address then responds to the DNS query. In the case where the nearest server becomes unavailable, the next nearest server responds to the query. OSPF anycast provides a dynamically routed failover to ensure that DNS can always resolve client requests within the AS. From the client perspective, anycasting is transparent and the group of DNS servers with the anycast address appears to be a single DNS server.

### Anycast Addressing for DNS Using OSPF



After you configure or change DNS anycast settings, you must restart the DNS services for the settings to take effect. When you enter any OSPF command and wait for the interface to return more information, the appliance disconnects your CLI session after you restart services or make other OSPF configuration changes through Grid Manager. Re-enter your credentials to log back in to the CLI. (For information, refer to the *Infoblox CLI Guide*.) To enable the appliance to support OSPF and advertising anycast addresses on OSPF from the loopback, you must first configure the LAN1 or LAN1 (VLAN) interface as an OSPF advertising interface. For information about VLAN, see [About Virtual LANs](#). You can also configure authentication for OSPF advertisements to ensure that the routing information received from a neighbor is authentic and the reachability information is accurate. This process can be implemented for OSPF over IPv4 networks but is not supported for IPv6/OSPFv3.



#### Note

For more information about the OSPF routing protocol, refer to *RFC 2328 "OSPFv2" and RFC 5340 "OSPF for IPv6"*.



## Configuring OSPF on the NIOS Appliance



### Note

Use the CLI command **show ospf** or `show ipv6_ospf` to display configuration and statistical information about the OSPF protocol running on the appliance. You can also use the **set ospf** or `set ipv6_ospf` command to write OSPF statistical information to the syslog. In the NIOS appliance, configuration of OSPF is limited to Syslog and the DNS anycast application.

To support DNS anycast and other routing-dependent applications on NIOS appliances, you must first configure the LAN1 or LAN1 (VLAN) interface as an OSPF advertising interface, and then assign an area ID on the interface to associate it with a specific OSPF area. The interface advertises the OSPF routing information to the network so that routers can determine the best server to query. Note that the appliance automatically uses the HA interface as the advertising interface for an HA pair, even though you select the **LAN1** interface. For anycasting, the advertising interface sends out routing advertisements about the anycast address into the network out to upstream routers.



### Note

IPv6 is not supported for the **Stub** and **Not-so-stubby** area types.

To configure the LAN1 (HA) or LAN1(VLAN) interface to be an OSPF advertising interface, perform the following tasks:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. Select the **Anycast** tab in the *Grid Member Properties* editor. The **Anycast Interfaces** appear in a table. You can add new anycast interfaces when needed.
3. Click the Add icon of the OSPF Area Configuration table and choose **IPv4 Configuration** or **IPv6 Configuration** to define a new OSPF Area. The OSPF Area Configuration will show a similar set of **Add (IPv4/IPv6) OSPF Area** configuration settings based on the protocol type. Enter the following information to configure the LAN1, or LAN1 (VLAN) as the OSPF advertising interface:
  - **Advertising Interface:** Displays the interface that sends out OSPF routing advertisement. OSPF advertisements are supported on the LAN1 and LAN1(VLAN) interfaces. For an HA pair, select **LAN1** and the appliance automatically uses the HA interface as the advertising interface.
  - **Area ID:** Enter the OSPF area identifier of the network containing the upstream routers, in either an IP address format or a decimal format. All network devices configured with the same OSPF area ID belong to the same OSPF area. The area ID configured on the Grid member must match the area ID of the upstream router configuration. Area ID numbers are defined in the same format for IPv6 and IPv4.
  - **Area Type:** Select the type of OSPF area to associate with the advertising interface from the drop-down list. The area type configured on the Grid member must match the area type of the upstream router configuration. The supported area types are described as follows:
    - **Standard:** A standard area has no restrictions on routing advertisements, and connects to the backbone area (Area 0) and accepts both internal and external link-state advertisements.
    - **Stub:** A stub area is an area that does not receive external routes.
    - **Not-so-stubby:** A not-so-stubby area (NSSA) imports autonomous system (AS) external routes and sends them to the backbone, but cannot receive AS external routes from the backbone or other areas.  
OSPF for IPv6 (known as OSPFv3) configuration does not support OSPF authentication options.
  - **AuthenticationType:** Select the authentication method to use to verify OSPF routing advertisements on the interface. The authentication type configured on the Grid member must match the authentication type of the upstream router configuration. The supported authentication types are described as follows:
    - **None:** No authentication for OSPF advertisement.
    - **Simple:** A simple password for OSPF advertisement authentication, in clear text.
    - **MD5:** An MD5 hash algorithm to authenticate OSPF advertisements. This is the most secure option.

- **Authentication Key ID:** Enter the key identifier to use to specify the correct hash algorithm after you select **MD** as your OSPF authentication type. The authentication key ID configured on the Grid member must match the authentication key ID of the upstream router configuration.
  - **Authentication Key:** Enter the authentication password to use to verify OSPF advertisements after you select **Simple** or **MD** as your OSPF authentication type. Specify a key string between 1 to 8 characters for Simple authentication, and a string between 1 to 16 characters for MD5 authentication. The authentication key configured on the Grid member must match the authentication key of the upstream router configuration.
  - **Cost:** Select one of the following:
    - **Calculate Automatically:** Select this checkbox to auto generate the cost to associate with the advertising OSPF interface to the appliance. If this checkbox is not selected, then you specify the cost value explicitly. Calculate the cost as 100,000,000 (reference bandwidth) divided by the interface bandwidth. For example, a 100Mb interface has a cost of 1, and a 10Mb interface has a cost of 10.
    - **Fixed Metric:** Enter the cost to associate with the advertising OSPF interface to the appliance.
  - **Hello Interval:** Specify how often to send OSPF hello advertisements out from the appliance interface, in seconds. Specify any number from 1 through 65,535. The default value is 10 seconds. The hello interval configured on the Grid member must match the hello interval of the upstream router configuration.
  - **Dead Interval:** Specify how long to wait before declaring that the NIOS appliance is unavailable and down, in seconds. Specify any number from 1 through 65,535. The default value is 40 seconds. The dead interval configured on the Grid member must match the dead interval of the upstream router configuration.
  - **Retransmit Interval:** Specify how long to wait before retransmitting OSPF advertisements from the interface, in seconds. Specify any number from 1 through 65,535. The default value is 5 seconds. The retransmit interval configured on the Grid member must match the retransmit interval of the upstream router configuration.
  - **Transmit Delay:** Specify how long to wait before sending an advertisement from the interface, in seconds. Specify any number from 1 through 65,535. The default value is 1 second. The transmit interval configured on the Grid member must match the transmit interval of the upstream router configuration.
  - Click **Add** to add the interface to the table.  
The **Cost**, **Hello Interval**, **Dead Interval**, **Retransmit Interval** and **Transmit Delay** settings can be configured for IPv6 deployments. OSPF authentication is not supported for IPv6 on the NIOS platform.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Managing OSPF

- OSPF advertises the route when the DNS service starts. The start DNS command creates an interface and starts the OSPF daemon.
- OSPF stops advertising the route when the DNS service stops. The stop DNS command stops the OSPF daemon and deletes the interface.
- The NIOS application does not support a route flap. For example, temporary DNS downtime such as restart, does not stop or re-instate the OSPF advertisement.
- The OSPF advertisement stops if DNS service is down for more than 40 seconds.

## Anycast and BGP4



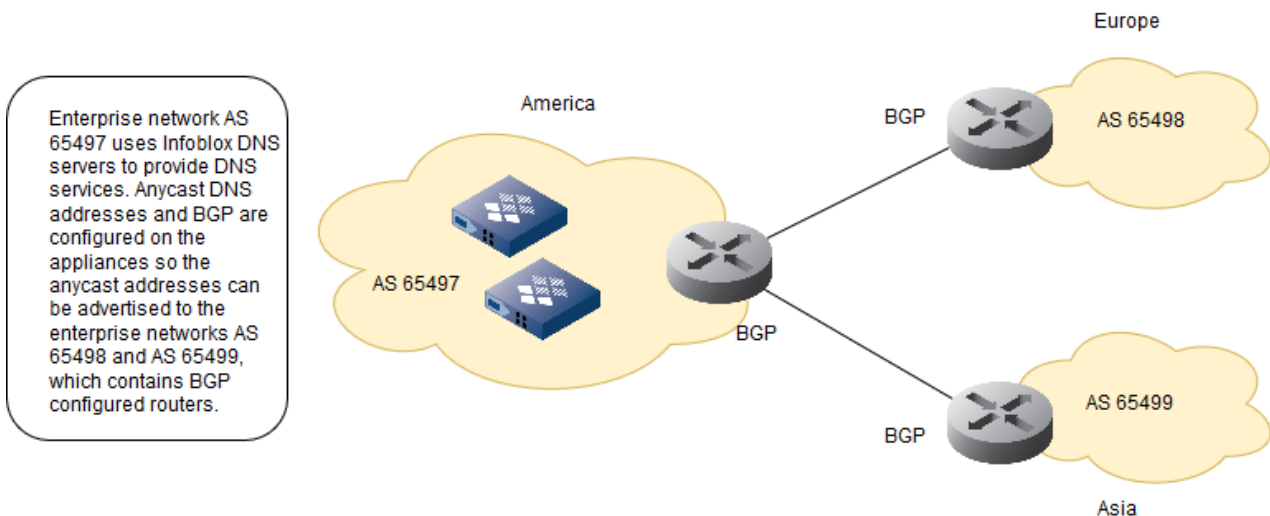
#### Note

Use the CLI command `show bgp` or `show ipv6_bgp` to display configuration and statistical information about the Border Gateway Protocol running on the appliance. You can also use the `set bgp` command to write OSPF statistical information to the syslog. In the NIOS appliance, configuration of BGP is limited to Syslog and the DNS anycast application.

BGP4 (henceforth referred to as BGP) is designed to distribute routing information among ASs and exchange routing and reachability information with other BGP systems using a destination-based forwarding paradigm. Unlike OSPF, which calculates routes within a single AS, BGP is a vector routing protocol that distributes routing information among different ASs. A unique ASN (autonomous system number) is allocated to each AS to identify the individual network in BGP routing. A BGP session between two BGP peers is an eBGP (external BGP) session if the BGP peers are in different ASs. A BGP session between two BGP peers is an iBGP (internal BGP) session if the BGP peers are in the same AS. BGP configuration enables large enterprises using BGP as the internetworking protocol to provide resilient DNS services using the Infoblox solution. While BGP is mostly used by ISPs, it is also used in larger enterprise environments that must interconnect networks that span geographical and administrative boundaries. In these environments, it is required to use BGP to advertise anycast routes. Using BGP allows the appliance to advertise DNS anycast addresses to neighboring routers across multiple ASs that also use BGP as their routing protocols.

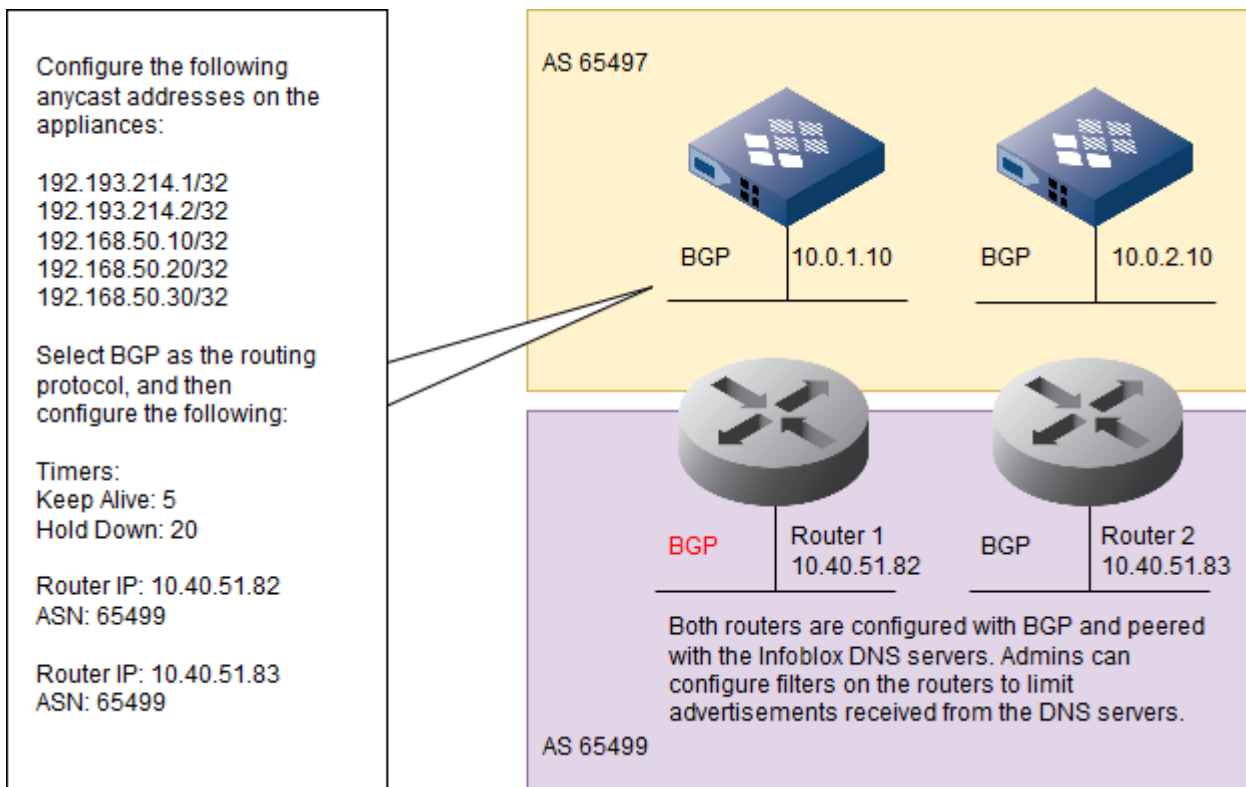
As illustrated in the figure *Anycast Addressing for DNS using BGP* below, to enable anycast for DNS queries among three different networks that span different geographical regions, you can configure two DNS servers with the same DNS anycast addresses in the AS 65497 network. Since other network routers in AS 65498 and AS 65499 also use BGP as the routing protocol, the DNS anycast addresses can be advertised across these networks.

#### *Anycast Addressing for DNS using BGP*



To enable DNS anycast addressing across different ASs, you configure BGP as the routing protocol on the NIOS appliance. (As illustrated in the figure *Anycast and BGP Configuration on Infoblox Appliances* below, the AS 65497 network contains the Infoblox DNS anycast servers, and the AS 65499 network contains Router 1 and 2. The routers use BGP and are peered with the DNS servers. You can configure anycast addressing on the loopback interface of the DNS servers and select BGP as the protocol to advertise the anycast addresses to Router 1 and 2 in AS 65499. For information, see [Configuring Anycast Addresses](#). Once you have configured the DNS servers, the appliances automatically add filters on the advertising interfaces to limit the advertisements to the configured anycast IP addresses. Similarly, BGP filters are applied to ensure that the DNS servers only receive default route advertisements from the neighboring routers.

#### *Anycast and BGP Configuration on Infoblox Appliances*



BGP uses timers to determine how often the appliance sends keepalive and update messages, and when to declare a neighboring router out of service. You can configure the time intervals for these timers. For information, see [Configuring BGP in the NIOS Appliance](#) below.

The BGP protocol service is automatically configured to send SNMP queries about BGP runtime data. The appliance sends SNMP traps to its neighboring routers when it encounters issues with the protocol. BGP is configured to send SNMP traps as defined in *RFC4273 Definitions of Managed Objects for BGP-4*. You must enable and configure the SNMP trap receiver on the Grid member for the member to send SNMP traps. For information, see [SNMP MIB Hierarchy](#).

You can use the `set bgp` command to set the verbosity levels of the BGP routing service. The appliance writes BGP statistical information to the syslog. After you configure the settings, you must restart the DNS services for the settings to take effect. For information, refer to the *Infoblox CLI Guide*. Note that when you enter any BGP command and wait for the interface to return more information, the appliance disconnects your CLI session if you restart services or make other BGP configuration changes through Grid Manager. You must re-enter your credentials to log back in to the CLI. You can configure BGP on any interface to advertise anycast addresses across multiple ASs.



Note

NIOS selects the interface for BGP advertisement based on the routing configuration.

The appliance supports BGP version 4. For more information about BGP, refer to *RFC4271, A Border Gateway Protocol 4 (BGP-4)*.

### Configuring BGP in the NIOS Appliance

You can configure the appliance as a BGP advertising interface for anycast addresses. The NIOS appliance advertises the BGP routing information to the network so routers can determine the nearest server to query. The NIOS appliance does not perform dynamic routing itself; it can use dynamic routing protocols to advertise its DNS anycast availability. You must define the ASN of the interface and list any neighboring routers that will receive the BGP announcements. On an HA pair, BGP runs only on the active node. In an HA failover, the BGP service resumes on the new active node.



#### Note

If you encounter Malformed AS\_PATH error, then remove the dont-capability-negotiate option. Infoblox doesn't provide an option to create confederation of autonomous systems if the BGP peer is configured by enabling the dont-capability-negotiate option.

## Authenticating BGP Neighbors

You can configure authentication for BGP advertisements to avoid any malicious interference by ASs. This ensures that the routing information exchanged between BGP peers is authentic, and it is accepted only if the authentication is successful. BGP authentication must be configured with the same password on both BGP peers. Otherwise, the connection between them is not established. The Infoblox BGP authentication fully conforms to RFC 2385. For information about BGP authentication, refer to RFC 2385, *Protection of BGP Sessions via the TCP MD5 Signature Option*.



#### Note

If you upgrade from a previous NIOS version to NIOS 6.11.0 or later, BGP authentication is disabled for existing BGP neighbors.

The BGP service restarts automatically when any of the following authentication changes are made:

- MD5 authentication is enabled or disabled for a BGP neighbor.
- Change the authentication password of a BGP neighbor, for which MD5 authentication is enabled.

To configure BGP for anycast addresses:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. In the *Grid Member Properties* editor, select the **Anycast** tab.
3. In the BGP Configuration section, complete the following:
  - **Interface Link Detection:** Select this checkbox to enable link detection when the default connection fails. This enables the router to track the next available route. For example, if LAN1 is set as the default gateway when both LAN1 and LAN2 are working, and LAN1 later fails, the router will switch to LAN2. When LAN1 reconnects, the router will then switch back to LAN1.
  - **ASN:** Enter the autonomous system number of the interface. You can enter an ASN number from 1 to 4294967295. You can configure only one ASN on each Grid member.
  - **BGP Timers:** BGP uses timers to control how often the interface sends KEEPALIVE messages and how long it waits before declaring a neighboring router out of service. The keepalive timer determines the time interval at which the interface sends KEEPALIVE messages to a neighboring router to inform the neighbor that the appliance is alive. The hold down timer determines how long the interface waits to hear a KEEPALIVE or UPDATE message before it assumes its neighbor is out of service. If a neighboring router is down, the interface terminates the BGP session and withdraws all the BGP routing information to the neighbor.
    - **Keep Alive:** Enter the time interval in seconds when the interface sends keepalive messages. You can enter a time from 1 to 21845 seconds. The default is four seconds.
    - **Hold Down:** Enter the time in seconds that the interface waits to hear a keepalive message from its neighbor before declaring the neighbor out of service. You can enter a time from 3 to 65535 seconds. The default is 16 seconds.Click the Add icon to add a neighboring router to receive BGP advertisements from the NIOS appliance. The appliance adds a new row to the table. Complete the following:
  - **Neighbor Router IP:** Enter the IP address (IPv4 or IPv6) of the neighboring BGP router. The neighboring router can be within the same AS (the most likely case) or from a router in an external AS.
  - **Remote ASN:** Enter the ASN of the neighboring router. You can enter an ASN number from 1 to 4294967295.
  - **MD5 Authentication:** Select this checkbox to enable MD5 authentication for the BGP neighbor. When you enable MD5 authentication, you must enter the authentication password in the **Password** field.

- **Password:** Enter the authentication password that the NIOS appliance uses to connect to the BGP neighbor. You can enter up to 80 printable ASCII characters. The password configured on the Grid member must match the password of the BGP neighbor. When you enter the password for a BGP neighbor, it will be preserved even if you disable MD5 authentication for the BGP neighbor later. But if you change the IP address for any existing BGP neighbor, you must re-enter the authentication password for the BGP neighbor, even if the authentication password remains the same.
  - **Comment:** Enter useful information about this neighboring router. Click the Add icon again to add another neighboring router. You can add up to 10 neighboring routers.
4. Save the configuration and click **Restart** if it appears at the top of the screen.
  5. Anycast configuration is complete. To activate anycast, see [Specifying Port Settings for DNS](#) and its subtopic, [Specifying Source Ports](#).

## About BFD (Bidirectional Forwarding Detection)

NIOS supports Anycast addressing for DNS using BGP and OSPF routing protocols. Since BGP and OSPF have timer granularity in seconds, the network re-convergence is slow in case of faults in forwarding path. BFD protocol is designed to provide faster failure detection using millisecond timer intervals. It can be enabled with routing protocols to achieve fast network re-convergence.

You can use BFD to detect failures early on and create adjacency with the next router. BFD can be enabled for OSPF or BGP and you can create BFD templates and assign it to OSPF Area or BGP neighbors. You can enable BFD simultaneously for OSPF and BGP, but only one BFD session will be created for a given neighbor. Infoblox recommends that you use the same BFD template for both the OSPF and BGP neighbors whenever such a configuration is required. When BFD is enabled on an anycast daemon, any failure on the DNS service also restarts the anycast daemon. When the anycast restart behavior is 'off', the BFD restarts only when there is a configuration change.

When the BFD protocol gets enabled, it by default also starts an internal DNS monitor to check whether the DNS service is up and running. The BFD DNS monitor works by sending a DNS query for the root zone "." to the configured anycast addresses. If the internal DNS monitor does not receive a DNS response, the DNS service is assumed to be down. The internal DNS monitor then sends the BFD down message and retracts the anycast route.

The BFD protocol feature is supported in NIOS 8.0 and later releases. This section provides a brief overview about enabling BFD for OSPF area and BGP neighbors, creating BFD templates, SNMP, and CLI commands.

### Warning

*The default advertised setting for BFD holddown is 300 ms (100 ms transit/receive intervals and detection multiplier 3). This setting is optimized for typical routers and directly connected endpoint configurations. If your network requires an implementation of L2 multi-path or port redundancy, you must adjust the holddown interval value higher than the spanning-tree rebalance latency to avoid unnecessary changes to the L3 network topology or the forwarding path for DNS traffic.*

### Note

You can enable or disable the **BFD Internal DNS Monitoring** checkbox only if you select the **Enable BFD** checkbox. When you enable the **BFD Internal DNS Monitoring** checkbox, Infoblox recommends that you also select the **Enable DNS Health Check** checkbox in the *Grid Properties Editor* or the *Member Properties Editor*. The **BFD Internal DNS Monitoring** checkbox is enabled by default.

## Enabling BFD for OSPF

You can enable BFD for both IPv4 and the IPv6 OSPF areas. To support DNS anycast and other routing-dependent applications on NIOS appliances, you must first configure the LAN1, LAN1 (VLAN), or HA (for HA pairs only) interface as an OSPF advertising interface, and then assign an area ID on the interface to associate it with a specific OSPF area.

To enable BFD for the IPv4 or the IPv6 OSPF area:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. In the *Grid Member Properties* editor, select the **Anycast** tab.
3. In the OSPF Area Configuration table, select the OSPF advertising interface, click the Edit icon, and then complete the following:
  - **Enable BFD**: Select this checkbox to enable BFD for the OSPF advertising interface.
  - **BFD template**: Click **Select BFD Template** and select a BFD template from the *Select BFD Template* dialog box. You can click **Clear** to remove the selected BFD template and select a new one.
  - **BFD Internal DNS Monitoring**: Select this checkbox to enable the internal DNS monitor to send and receive DNS responses and to retract the OSPF route if it does not receive a DNS response.
4. Save the configuration.

When OSPF session with a neighbor router in the OSPF Area reaches FULL state, BFD session is automatically created. By default, BFD runs with no authentication and timer intervals of 100ms transmit, 100ms receive and multiplier 3 (hold down time = 300ms). The actual runtime intervals are negotiated with the peer as per BFD standard RFC 5880. If these intervals are not suitable or authentication needs to be enabled for BFD, you must create a BFD template.

### Enabling BFD for OSPF

The screenshot shows the 'Basic' tab of the 'Edit IPv4 OSPF Area' configuration. The 'Advertising Interface' is set to 'LAN1'. The 'Area ID' is '34' and the 'Area Type' is 'Standard'. Under 'Authentication', the 'Type' is 'None' and the 'Key ID' is '1'. The 'Cost' section has 'Calculate Automatically' selected and 'Fixed Metric' set to '5'. The 'Hello Interval(s)' is '10', 'Dead Interval(s)' is '40', and 'Retransmit Interval(s)' is '5'. The 'Enable BFD' checkbox is checked. The 'BFD Template' dropdown is set to 'bfd\_template1' with a 'Select' button.

You can use the `show ipv6_ospf neighbor` CLI command to view runtime BFD information for OSPF.

### Enabling BFD for BGP Neighbor

BFD can be enabled for each configured BGP neighbor individually. You can also use the Enable Multi-hop option, which allows BGP to connect to BGP neighbors which are more than one IP hops away.

To enable BFD for the BGP neighboring router:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. In the *Grid Member Properties* editor, select the **Anycast** tab.
3. In the BGP Neighbor Configuration table, select the BGP neighboring router, click the Edit icon, and then complete the following:
  - **Enable Multi-hop (optional)**: Select this checkbox to allow BGP to connect with the neighbors which are more than one IP hops away.
  - **Hop Limit**: Enter the maximum hop limit. The default value is 255.



- **Enable BFD:** Select this checkbox to enable BFD for the BGP neighboring router.
- **BFD template:** You can assign a BFD template to the BGP neighboring router to run BFD with non-default settings. Click **Select BFD Template** and select a BFD template from the *Select BFD Template* dialog box. You can click **Clear** to remove the selected BFD template and select a new one.
- **BFD Internal DNS Monitoring:** Select this checkbox to enable the internal DNS monitor to send and receive DNS responses and to retract the BGP route if it does not receive a DNS response.

4. Save the configuration.

The BFD session for a given BGP neighbor is created when the BGP state reaches 'Established'.

### Enabling BFD for BGP Neighbor

### Viewing Runtime BFD Information for BGP

You can use the `show bgp neighbor` CLI command to view runtime BFD information for BGP. For more information, see [show bgp](#).

### Creating a BFD Template

BFD advertises the default hold-down interval of 300ms and authentication is disabled, by default. In order to configure faster or slower hold-down intervals, you can create BFD templates and assign it to the OSPF Area or BGP neighbors. You can configure a BFD template at the Grid level and assign it to multiple Grid members. The BFD template can be assigned to the BGP neighbor or OSPF Area of any Grid member in the Grid and it can be assigned to multiple BGP neighbors or OSPF Areas.

To create BFD templates:

1. From the **Grid** tab -> **Grid Manager** tab, expand the Toolbar and click **Manage BFD Templates**.
2. In the *Manage BFD Templates* wizard, click the Add icon, and then complete the following:
  - **Name:** Enter the name of the BFD template.
  - **Authentication Type:** Select the authentication type from the drop-down list. You can select one of the following authentication types: **MD5**, **SHA-1**, **Meticulous MD5**, or **Meticulous SHA-1**. The BFD authentication type fully conforms to RFC 5883.
  - **Authentication Key ID:** Enter the key identifier to use to specify the correct hash algorithm after you select the authentication type. If you do not enter a value here, the appliance by default sets 'one' as the

authentication key ID. The authentication key ID configured on the Grid member must match the authentication key ID of the upstream router configuration.

- **Authentication Secret/Password:** Enter the authentication password to use to verify after you select the authentication type. You can enter password with 4-16 printable ASCII characters. The authentication password configured on the Grid member must match the authentication key of the upstream router configuration.
- **Intervals:** Specify the following BFD timer intervals for each router interface.
  - **Min Rx Interval (ms):** Enter the minimum receive interval. The default is 100ms. The minimum receive interval value must be an integer between 50 and 9999 (both inclusive).
  - **Min Tx Interval (ms):** Enter the minimum transmit interval. The default is 100ms. The minimum transmit interval value must be an integer between 50 and 9999 (both inclusive).
  - **Multiplier:** Enter the detection multiplier. You can enter a value between 3 and 50. The default is 3.

3. Click **Add**.

After you have added BFD templates, you can do the following:

- Select a BFD template and click the Edit icon to edit the configuration.
- Select a BFD template and click the Delete icon to delete it.

### Manage BFD Templates

The screenshot shows the Infoblox Grid Administration interface. The main content area displays a table of BFD templates. The table has columns for NAME, HA, STATUS, IPV4 ADDRESS, IPV6 ADDRESS, IDENTIFY, DHCP, DNS, and TFTP. The status of each template is shown in a green box labeled 'Running'.

NAME	HA	STATUS	IPV4 ADDRESS	IPV6 ADDRESS	IDENTIFY	DHCP	DNS	TFTP
infoblox.localdc	No	Running	10.35.134.1		Unsupported			
ptmember.com	No	Running	10.35.0.44		Off			
discoveryment	No	Running	10.35.1.44		Off			
reportingmemb	No	Running	10.35.135.1		Unsupported			

## Manage BFD Templates Wizard

NAME	AUTHENTICATION ...	AUTHENTICATION ...	MIN RX INTERVAL	MIN TX INTERVAL	MULTIPLIER
bfd_template1	None	1	100	100	3

## Enabling and Disabling DNS Health Check Monitor

In order to minimize downtime for DNS and ensure high availability, NIOS implements DNS process monitoring and self-recovery on each Grid member, in order to minimize downtime for DNS and ensure high availability. You can enable the DNS health check monitor to monitor whether the DNS server is responding to client requests. When you enable this feature, the appliance sends a query to the DNS server and waits for the response until the specified timeout duration. If the appliance is unable to receive a response from the DNS server after the specified number of retries, the appliance sends SNMP traps and email notifications about the failure. The appliance performs the DNS health check periodically based on the specified time interval.

If BFD is used for anycast fault detection, the BFD session state advertised from the member can be in the Down state whenever there is a DNS health check failure. This allows quick anycast route tear-down and the network might converge with another DNS server that can serve same anycast IP.

Additionally, you can also configure domain names in the DNS health check monitor, which are probed simultaneously and if any one of the domains fail to resolve for consecutive attempts, the DNS health is considered as Down. If recursion is enabled on the Grid member, the queries to these domains help to assert the ability of the DNS server to reach the external authoritative servers and optionally trigger network re-convergence in case of a failure. When no domains are configured, local PTR queries are used to probe the DNS process.

### Warning

*The DNS health check monitor might not work properly if the DNS blackhole feature is enabled or if any named ACL is blocking the query sent to the loopback interface.*

To enable or disable the DNS health check monitor:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, and then select **Grid DNS Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.

In the *Grid DNS Properties* or the *Member DNS Properties* editor, click **Toggle Advanced Mode** if the editor is in the basic mode.

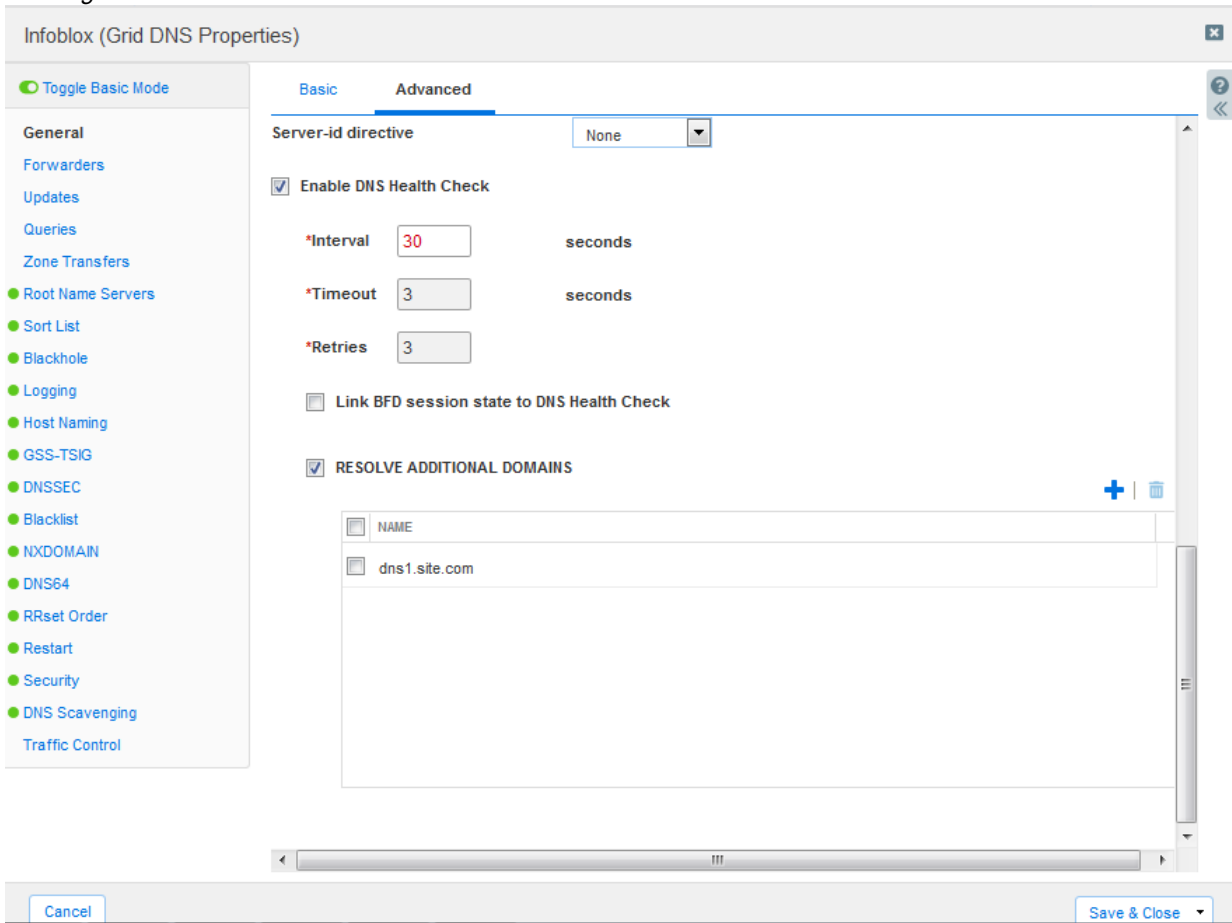
2. Click the **Advanced** subtab of the **General** tab and then complete the following:
  - **Enable DNS Health Check:** This checkbox is deselected by default, meaning the DNS health check monitor is disabled. Select this checkbox to enable the DNS health check monitor and specify the following:
    - **Interval:** Enter the time interval in seconds. The interval value is measured from the end of the previous monitoring cycle. The default is 30 seconds. You can enter a value between 10 and 21600 seconds.
    - **Timeout:** Enter the timeout value in seconds. This is the time the appliance waits for a response to the query. The default is 3 seconds. You can enter a value between 1 and 10 seconds.
    - **Retries:** Enter the number of times the appliance tries to send the query after a failed attempt. The default is 3. You can enter a value between 1 and 10.
    - **Link BFD session state to DNS Health Check:** Select this checkbox to link the BFD session state with the DNS health check monitor.
    - **Resolve Additional Domains:** Click the Add icon and enter the domain name. The DNS health check monitor sends recursive queries to the local DNS server (BIND/Unbound) for the domain names listed in this table. You can add up to 16 domain names.
3. Save the configuration.



#### Note

You must carefully select the domain names for DNS health check monitor with BFD session in order to avoid unnecessary changes in downstream DNS traffic due to transient health check query failures. Setting a higher timeout or retry count might help in avoiding false alarms.

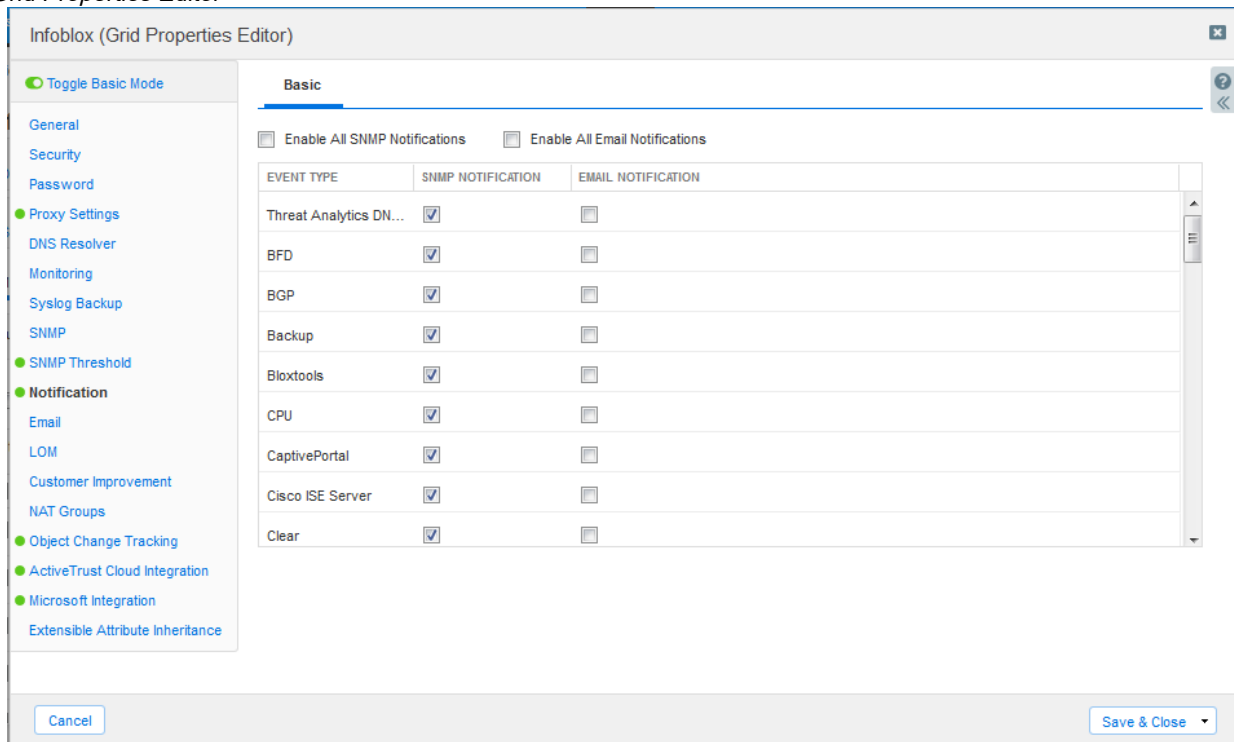
### Enabling DNS Health Check Monitor



### Monitoring with SNMP

Infoblox MIBs (IB-TRAP-MIB, IB-PLATFORMONE-MIB) are updated to include a notification for BFD process failure (ibBFDSOFTWAREFAILURE). By default, SNMP notifications are enabled for the BFD process failure event. You can enable or disable SNMP and email notifications for specific event types, by selecting the corresponding checkboxes in the **Notification** tab of the *Grid Properties* or *Member Properties* editor.

## Grid Properties Editor



In addition, the BFD process can generate SNMP traps for session state changes according to the standard BFD MIBs described in *RFC 7330* and *RFC 7331*:

- .1.3.6.1.2.1.222.0.1 (bfdSessUp): This notification (aka trap) is sent when one of the neighbors changes the BFD-session state as 'Up.'
- .1.3.6.1.2.1.222.0.2 (bfdSessDown): This notification (aka trap) is sent when one of the neighbors changes the BFD-session state as 'Down' or 'AdminDown.'
- .1.3.6.1.2.1.222.1.2.1.13 (bfdSessDiag): The diagnostic code which can be one of the following:
  - noDiagnostic (0)
  - controlDetectionTimeExpired (1)
  - echoFunctionFailed (2)
  - neighborSignaledSessionDown (3)
  - forwardingPlaneReset (4)
  - pathDown (5)
  - concatenatedPathDown (6)
  - administrativelyDown (7)
  - reverseConcatenatedPathDown (8)
  - misConnectivityDefect (9)

Note that you must download the following MIBs to enable the trap-receiver to parse the notifications:

- BFD-STD-MIB
- BFD-TC-STD-MIB
- DIFFSERV-MIB
- DIFFSERV-DSCP-TC
- INTEGRATED-SERVICES-MIB
- IANA-BFD-TC-STD-MIB

## DHCP

This section describes how to configure the Grid to provide DHCP services. It includes the following topics:

- [Configuring Infoblox DHCP Services](#)
- [Configuring DHCP Properties](#)
- [Managing IPv4 DHCP Data](#)
- [Managing DHCP Templates](#)
- [Managing IPv6 DHCP Data](#)
- [Configuring DHCP Failover](#)
- [Configuring DHCP Filters](#)
- [Authenticated DHCP](#)
- [Managing Leases](#)

## Configuring Infoblox DHCP Services

DHCP (Dynamic Host Configuration Protocol) is a network application protocol that automates the assignment of IP addresses and network parameters to DHCP-configured network devices (DHCP clients). When a DHCP client connects to a network, it sends a request to obtain an IP address and configuration information from the DHCP server. The DHCP server manages a pool of IP addresses and configuration information such as default gateway, domain name, and DNS server. Depending on the configuration, the DHCP server either assigns or denies an IP address to a client request. It also sends network configuration parameters to the client.

You can configure a NIOS appliance to provide DHCP service for IPv4 and IPv6. The Infoblox DHCP server complies with a number of DHCP and DHCPv6 RFCs (see Appendix A Product Compliance). Limited-access admin groups can access certain DHCP resources only if their administrative permissions are defined. For information on setting permissions for admin groups, see [Managing Administrators](#).

This section provides an overview of the Infoblox DHCP services for IPv4 and IPv6. It contains the following topics:

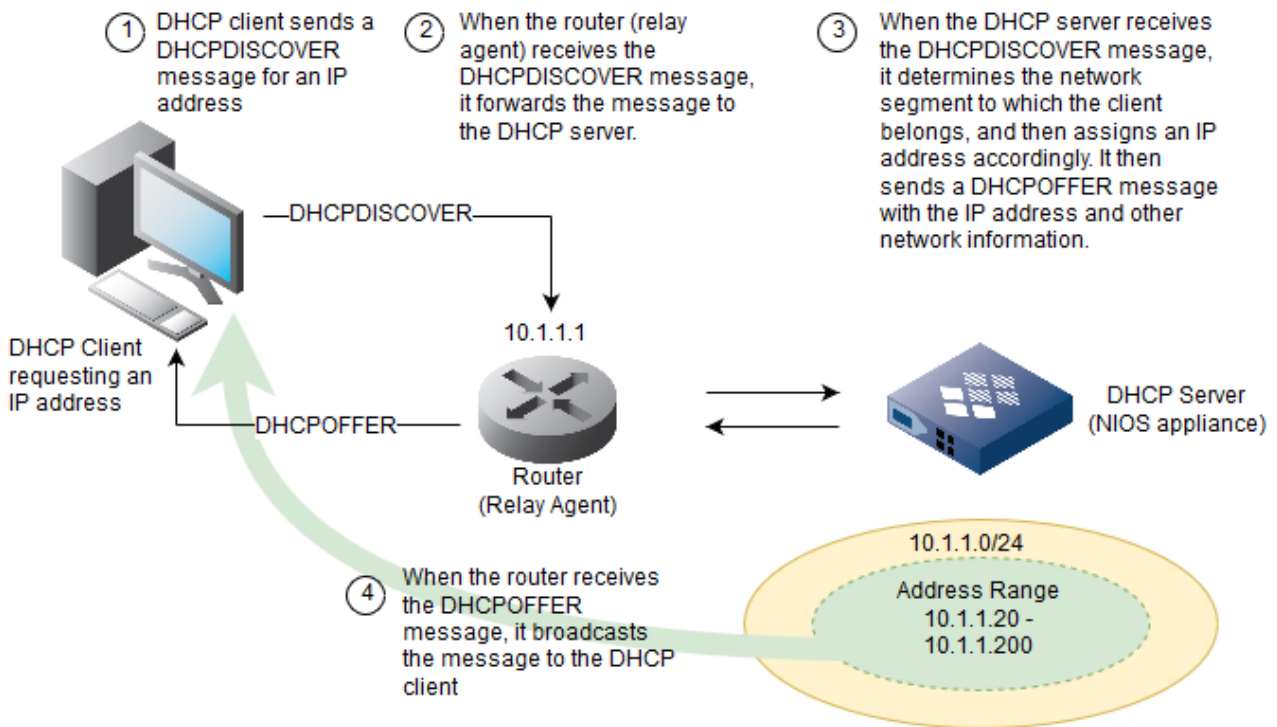
- [IPv4 DHCP Protocol Overview](#)
- [IPv6 DHCP Protocol Overview](#)
- [Configuring DHCP Overview](#)
- [Managing DHCP Data](#)
- [DHCP Inheritance](#)
- [Configuring Network Views](#)

### IPv4 DHCP Protocol Overview

As illustrated in the figure below, when a DHCP client requests an IP address, it sends a DHCPDISCOVER message to the router, which can act as a relay agent. The router forwards the message to the DHCP server. When the DHCP server receives the DHCPDISCOVER message, it determines the network segment to which the client belongs and assigns an IP address. The DHCP server then sends a DHCPOFFER message that includes the IP address and other network configuration information. When the router receives the DHCPOFFER message, it broadcasts the message to the client that sent the DHCPDISCOVER message.

*IP Address Allocation Process*



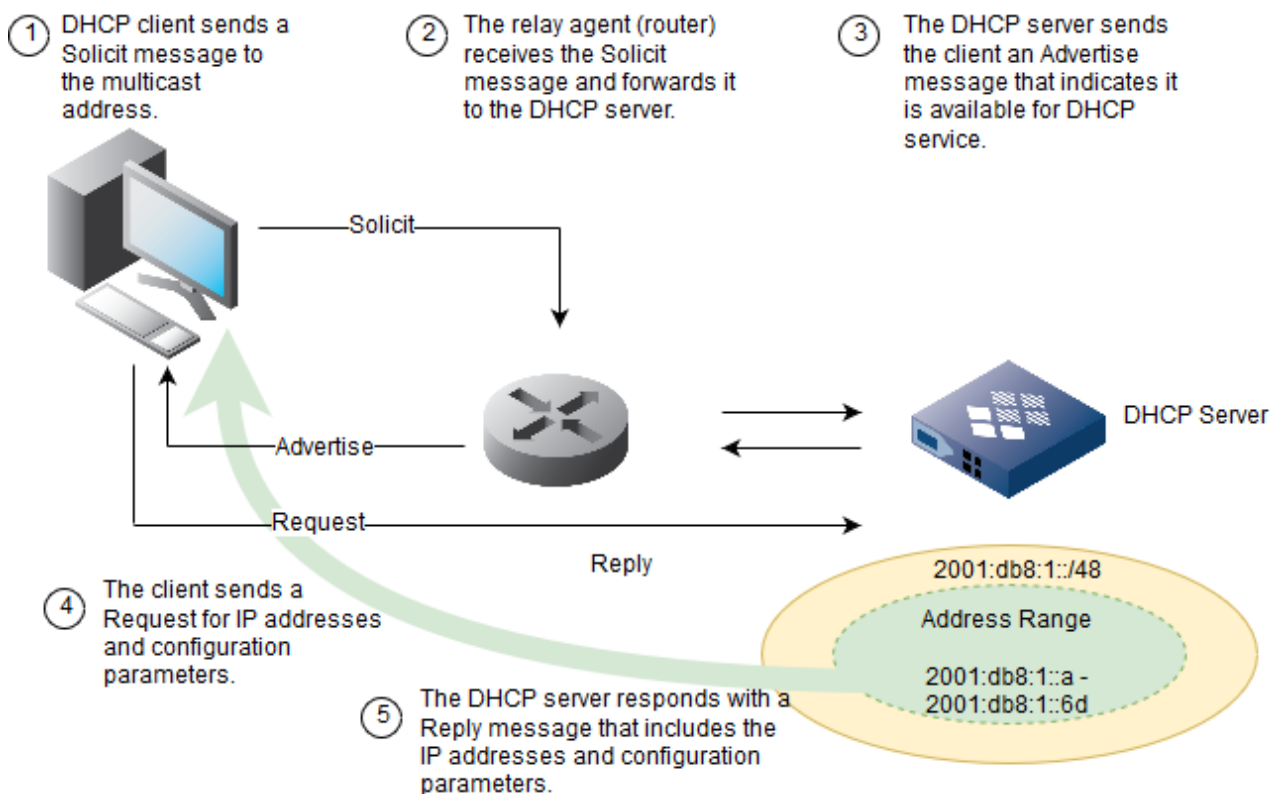


## IPv6 DHCP Protocol Overview

You can configure NIOS appliances to support DHCP for IPv6 (DHCPv6), the protocol for providing DHCP services for IPv6 networks. The DHCPv6 client-server model is similar to that of IPv4. DHCP clients and servers use a reserved, link-scoped multicast address to exchange DHCP messages. When a DHCP client needs to send messages to a DHCP server that is not attached to the same link, a DHCP relay agent can be used to relay messages between the client and server. Each IPv6 DHCP server and client has a unique DHCP unique identifier (DUID). DHCP servers use DUIDs to identify clients when providing configuration parameters, and clients use DUIDs to identify the source of the DHCP messages from servers.

As illustrated in the figure Client DHCPv6 Configuration Workflow below, a DHCP client that needs an IPv6 address sends a Solicit message to the well-known multicast address. DHCPv6 servers then send Advertise messages to the client to indicate that they are available. The client sends a Request message to a specific DHCPv6 server to request IP addresses and configuration parameters. The DHCPv6 server responds with a Reply message that contains the IP addresses and configuration parameters. You can view statistics about the IPv6 messages on the Dashboard.

*Client DHCPv6 Configuration Workflow*




Infoblox DHCP servers also supports stateless configuration in which a DHCP client does not need IP addresses, but needs configuration information only. The DHCP client sends an Information-Request packet to obtain configuration information, and the server sends a Reply message with the requested information. For more information, refer to *RFC 2462, IPv6 Stateless Address Autoconfiguration*.

### IPv6 Address Structure

An IPv6 address consists of the following:

- Global Routing Prefix — Global routing prefix is a (typically hierarchically-structured) value assigned to a site. For example, an ISP can delegate a prefix to your site, which you can then divide into subnets.
- Subnet ID — Subnet ID is an identifier of a link within the site.
- Interface ID — Interface Identifier. This portion of the address identifies the interface on the subnet. This is equivalent to the host identifier for IPv4 addresses.

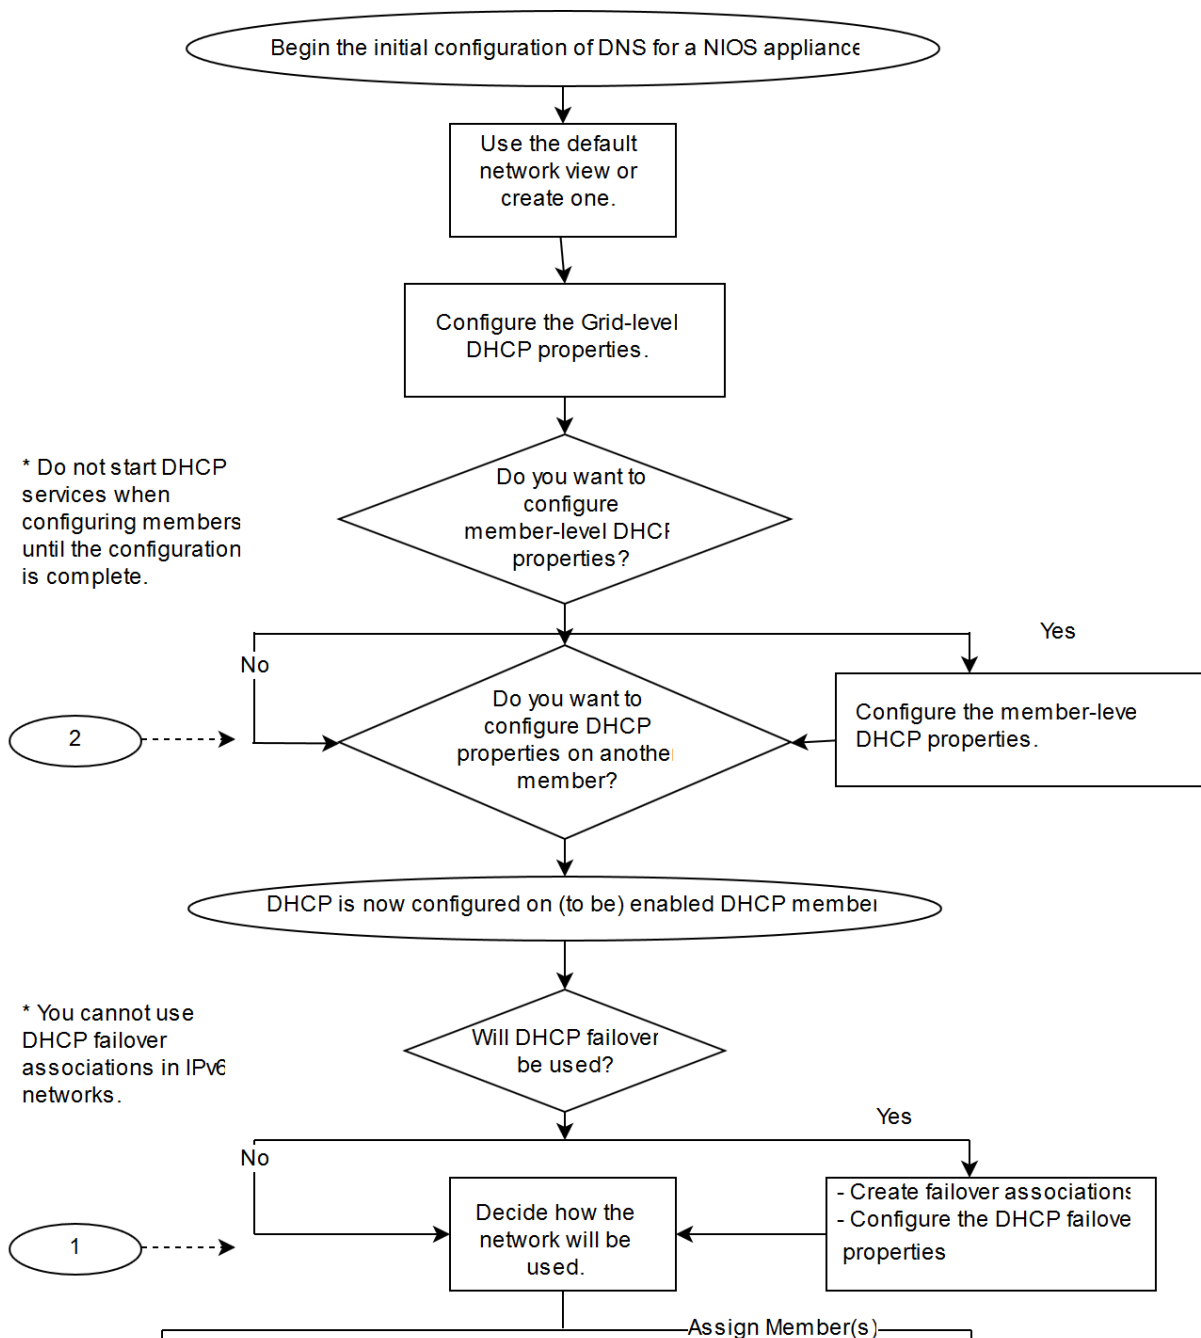
#### IPv6 Address Structure

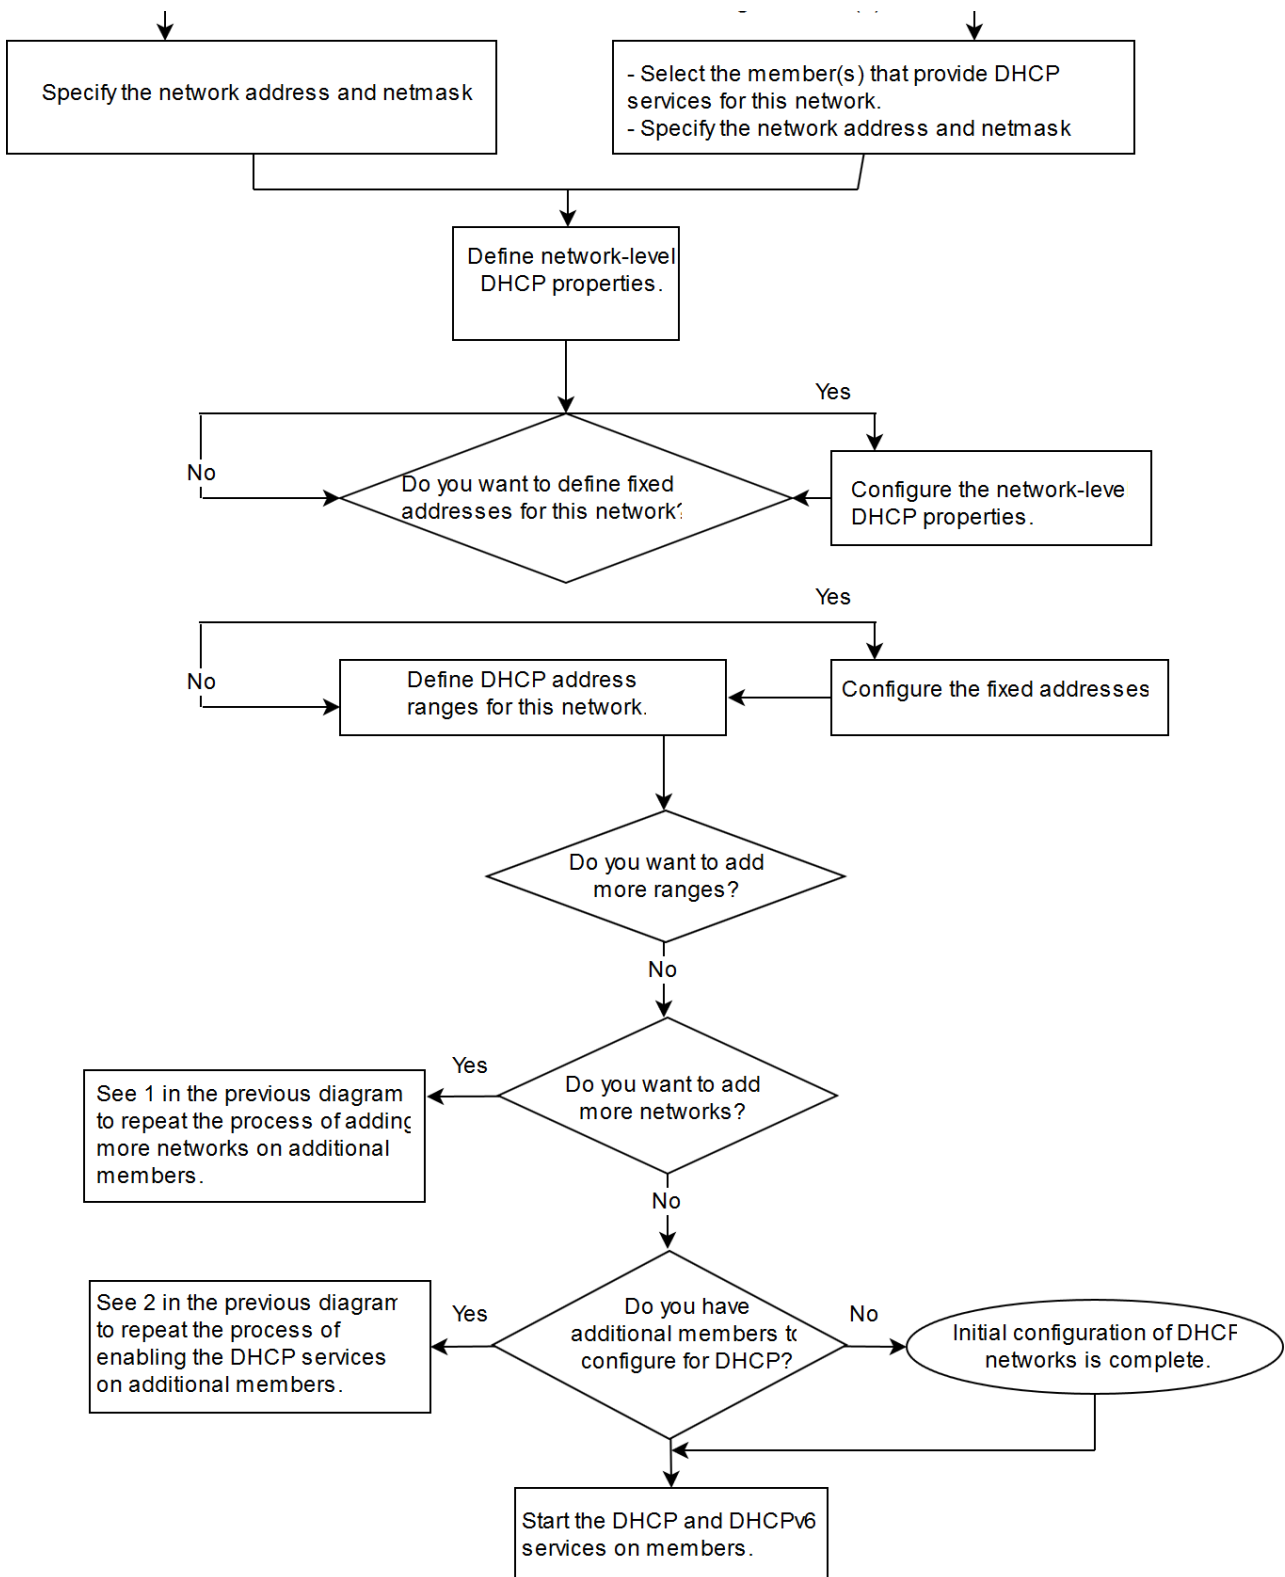


When you enter an IPv6 address in Grid Manager, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2006:0000:0000:0123:4567:89ab:0000:cdef can be shortened to 2006::123:4567:89ab:0:cdef. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The NIOS appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered.

## Configuring DHCP Overview

An overview of the complete DHCP configuration process is outlined in the following diagram, illustrating the main steps for preparing a NIOS appliance for use. Note that the process for configuring the DHCP server is the same for IPv4 and IPv6 networks, except that failover associations are not supported in IPv6 networks.





## Configuring the Connecting Switch

To ensure that VRRP (Virtual Router Redundancy Protocol) works properly, configure the following settings at the port level for all the connecting switch ports (HA, LAN1, and LAN2):

- Spanning Tree Protocol: Disable.  
For vendor specific information, search for "HA" in the Infoblox Knowledge Base system at <https://support.infoblox.com>.
- Trunking: Disable.  
If VLAN tagging is enabled on an Infoblox HA appliance, you must enable trunking at the port level.
- EtherChannel: Disable.
- IGMP Snooping: Disable.
- DHCP Snooping: Disable or Enable Trust Interface.  
You must disable DHCP Snooping to successfully run DHCP services on the Grid. For more information about DHCP services, see [Configuring Infoblox DHCP Services](#).
- Port Channeling: Disable.
- Speed and Duplex settings: Match these settings on both the Infoblox appliance and switch.
- Disable other dynamic and proprietary protocols that might interrupt the forwarding of packets.  
By default, a NIOS appliance automatically negotiates the optimal connection speed and transmission type (full or half duplex) on the physical links between its LAN1 or LAN1 (VLAN), HA, and MGMT ports and the Ethernet ports on the connecting switch. If the two appliances fail to auto-negotiate the optimal settings, see as described in [Modifying Ethernet Port Settings](#) for steps you can take to resolve the problem.

## Managing DHCP Data

You can configure a NIOS appliance to provide DHCP service for IPv4 and IPv6, and manage both IPv4 and IPv6 objects. When you define DHCP objects, you can track specific information about a network device by defining extensible attributes. Extensible attributes are fields that you define to track properties such as network locations or device models. For more information, see [Managing Extensible Attributes](#).

### About Networks

You can configure DHCP IPv4 and IPv6 properties for the Grid and its members, and then define the IPv4 and IPv6 networks that they serve.

All networks, both IPv4 and IPv6, must belong to a network view. The appliance has one default network view and unless you create additional network views, all networks belong to the default view. Note that because network views are mutually exclusive, you can create networks with overlapping IP address spaces in two different network views. For more information, see [Configuring Network Views](#).

#### Note

The 255.255.255.255 limited broadcast address is reserved. The appliance does not automatically create glue A records for this address. You can however create an NS record without the associated glue records. For more information, see [Changing the Interface IP Address](#).

### About Shared Networks

A shared network is a network segment to which you assign two or more subnets. When subnets in a shared network contain IP addresses that are available for dynamic allocation, the addresses are put into a common pool for allocation when client requests arise. When you create a shared network, the DHCP server can assign IP addresses to client requests from any subnet (that resides on the same network interface) in the shared network. For example, when you have networks A, B, and C on the same network interface and you assign them to a shared network, the DHCP server can allocate available IP addresses from any DHCP range within networks A, B, and C even when all the client requests originate from network A. When adding subnets to a shared network, ensure that the subnets are assigned to the same

members to avoid DHCP inconsistencies.

Before creating a shared network, you must first create the subnets. For example, you must first create the IPv4 networks `10.32.1.0` and `10.30.0.0` before designating them to a shared network or create the IPv6 networks `2001:db8:1::/48` and `2001:db8:2::/48` before designating them to a shared network.

After you create a network, you can define their DHCP resources such as DHCP ranges, fixed addresses, reservations, host records, and roaming hosts. IPv4 and IPv6 support most of the same DHCP objects, except that IPv6 does not support reservations.

## About DHCP Ranges

A DHCP range is a pool of IP addresses from which the appliance allocates IP addresses. You must add a DHCP address range in your network so the appliance can assign IP addresses to DHCP clients within the specified range. IPv6 DHCP ranges can also contain a range of IPv6 prefixes that it delegates to DHCP clients that request them. You must assign a DHCP range to a Grid member. Note that you can only assign DHCP ranges to members and networks that are in the same network view. If the server is an independent appliance, you must specify this appliance as the member that serves the DHCP range. In addition, you can also assign IPv4 DHCP ranges to failover associations.

## About Exclusion Ranges

You can define an exclusion range within a DHCP range. Creating an exclusion range prevents the appliance from assigning the addresses in the exclusion range to clients. IP addresses in an exclusion range are excluded from the pool of IP addresses. You can use exclusions to split a DHCP range into multiple blocks of ranges. You can also use addresses in the exclusion ranges as static IP addresses for network devices such as legacy printers that do not support DHCP. An exclusion in a range can help prevent address conflicts between statically configured devices and dynamically configured devices.

## About Fixed Addresses

You can configure fixed addresses for network devices, such as routers and printers, that are not frequently moved from network to network. By creating fixed addresses for network devices, clients can reliably reach them by their domain names. Some network devices, such as web or FTP servers, can benefit from having fixed addresses for this reason. In IPv4 and IPv6 networks, you can also reserve an IP address that is not part of a DHCP range by defining a reservation. For information about creating reservations, see [Configuring IPv4 Reservations](#).

## About Hosts

Infoblox hosts are data objects that contain DNS, DHCP, and IPAM data of the assigned addresses. You can assign multiple IPv4 and IPv6 addresses to a host. When you create a host, you are specifying the name-to-address and address-to-name mappings for the IP addresses that you assign to the host. For information about Infoblox hosts, see [About Host Records](#).

## DHCP Configuration Checklists

After you complete the appliance configuration for each member in the Grid, as described in [Managing Appliance Operations](#), you can configure DHCP services.

The following checklist includes the major steps for configuring DHCP service for IPv4:

### *IPv4 DHCP Configuration Checklist*

Step	For more information
Configure DHCP properties for the Grid and members.	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPv4 DHCP Properties</a></li> <li>• <a href="#">Understanding DDNS Updates from DHCP</a></li> <li>• <a href="#">Configuring DHCP IPv4 and IPv6 Common Properties</a></li> <li>• <a href="#">Configuring the Lease Logging Member</a></li> </ul>
Decide if you want to configure a DHCP failover association.	<ul style="list-style-type: none"> <li>• <a href="#">Configuring Failover Associations</a></li> </ul>
Configure networks based on your network requirements and decide if you want to override the Grid or member DHCP configuration for the networks	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPv4 Networks</a></li> <li>• <a href="#">Configuring IPv4 Shared Networks</a></li> </ul>
Decide if you want to configure fixed addresses and reservations, and whether to override the <a href="#">Configuring IPv4 Reservations</a> upper level DHCP properties for the fixed addresses and reservations.	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPv4 Fixed Addresses</a></li> <li>• <a href="#">Configuring IPv4 Reservations</a></li> </ul>
Define DHCP ranges and decide whether to override the upper level DHCP properties for the ranges.	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPv4 Address Ranges</a></li> </ul>
Enable DHCP services on the member.	<ul style="list-style-type: none"> <li>• <a href="#">Starting DHCP Services on a Member</a></li> </ul>

The following checklist includes the major steps for configuring DHCP service for IPv6:

#### *IPv6 DHCP Configuration Checklist*

Step	For more information
Configure DHCP properties for the Grid and members.	<ul style="list-style-type: none"> <li>• <a href="#">Configuring DHCPv6 Properties</a></li> <li>• <a href="#">Understanding DDNS Updates from DHCP</a></li> <li>• <a href="#">Configuring DHCP IPv4 and IPv6 Common Properties</a></li> <li>• <a href="#">Configuring the Lease Logging Member</a></li> </ul>
Configure networks based on your network requirements and decide if you want to override the Grid or member DHCP configuration for the networks.	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPv6 Networks</a></li> <li>• <a href="#">Configuring IPv6 Shared Networks</a></li> </ul>
Decide if you want to configure fixed addresses and reservations, and whether to override the upper level DHCP properties for the fixed addresses and reservations.	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPv6 Fixed Addresses</a></li> </ul>
Define DHCP ranges and decide whether to override the upper level DHCP properties for the ranges	<ul style="list-style-type: none"> <li>• <a href="#">Configuring IPv6 Address Ranges</a></li> </ul>
Enable DHCP services on the member.	<ul style="list-style-type: none"> <li>• <a href="#">Starting DHCP Services on a Member</a></li> </ul>

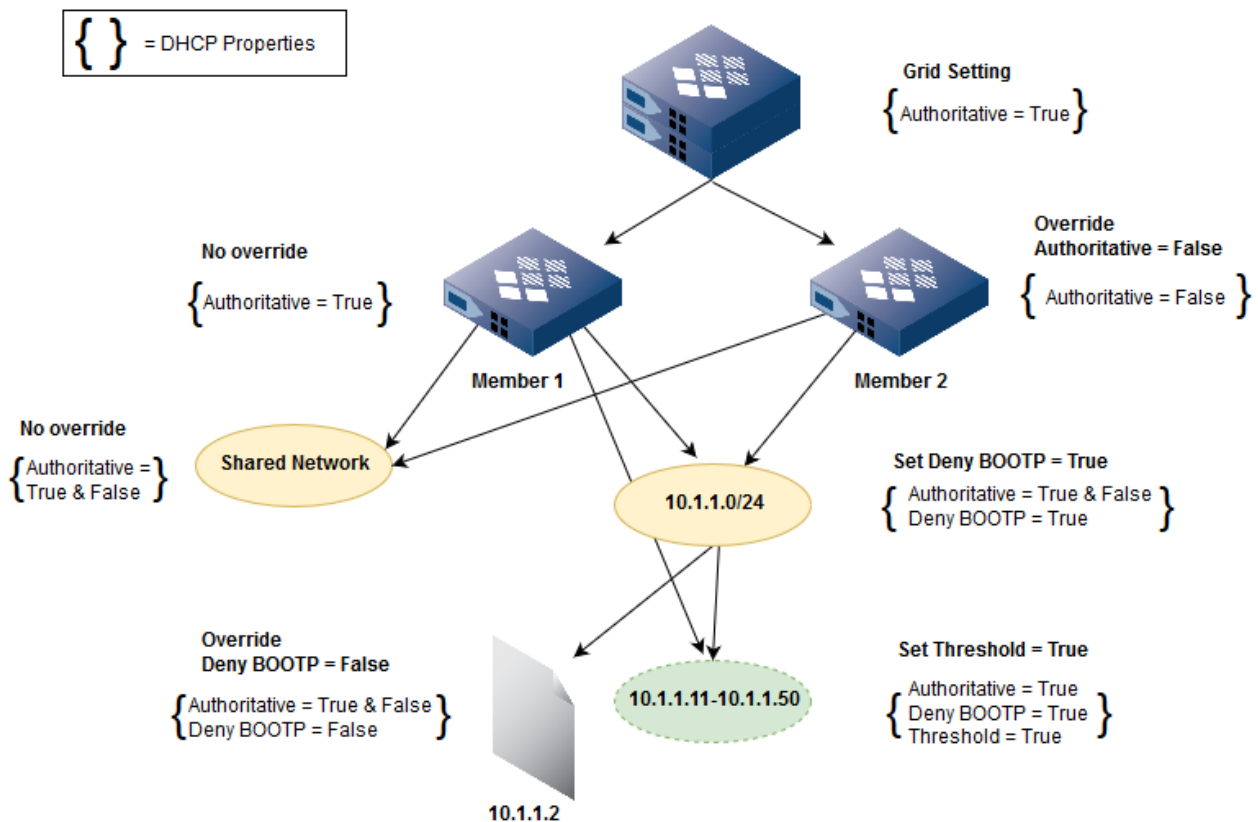


## DHCP Inheritance

When you configure DHCP properties for the Grid, members, networks, shared networks, DHCP ranges, fixed addresses, reservations, host addresses, and roaming hosts, the appliance applies the configured properties hierarchically. In addition, IPv4 DHCP objects inherit IPv4 specific properties and IPv6 objects inherit IPv6 specific properties. For example, when you set DHCP IPv4 properties for the Grid, all DHCP IPv4 objects inherit the properties from the Grid unless you override them at a specific level, and the same applies for IPv6 properties and objects. Properties set at the member level override Grid-level settings and apply to the objects that the member serves. Properties set at the network level override member-level settings and apply to the objects within the network. Properties set for a DHCP range override those set at higher levels. You can also set specific properties that apply only to fixed addresses, reservations, host addresses, and roaming hosts.

The figure [Inheritance Hierarchy in a Grid](#) below illustrates some inheritance scenarios that can occur in a Grid. As shown in the figure, the authoritative server configuration set for the Grid is inherited by the members. Since Member 1 has no overrides and Member 2 overrides the authoritative server configuration, they have different DHCP configurations. Grid Manager applies DHCP properties hierarchically from the Grid down. Therefore, a DHCP object below the member level can inherit DHCP properties with multiple values from multiple sources. In Inheritance Hierarchy in a Grid, network 10.1.1.0/24 inherits multiple values (True and False) from the members for the authoritative server configuration. The shared network, which includes 10.1.1.0/24, inherits DHCP properties from both members. For DHCP range 10.1.1.11 - 10.1.1.50, since Member 1 is the assigned member, it inherits properties from Member 1 and the network. The fixed address 10.1.1.2 overrides the BOOTP settings and inherits the authoritative server configuration from both members and the network.

*Inheritance Hierarchy in a Grid*



When a DHCP property contains inherited values from different sources, the appliance displays the corresponding information when you create or modify an object. Based on the information provided, you can then decide whether to

override or keep the inherited values. You must have read/write permissions to the DHCP resources to override inherited values. You can only view inherited values and paths if you have read-only permissions.

## Overriding DHCP Properties

DHCP properties configured at the Grid level apply to the entire Grid. You can choose to keep the inherited properties or override them when you configure the properties for a member, network, shared network, DHCP range, fixed address, host address, or roaming host. For example, you can override the values of DHCP properties inherited from a member and enter unique values for a network that is configured for DHCP.

To override an inherited value:

1. In a wizard or editor, click **Override** next to a property to enable the configuration. The **Override** button changes to **Inherit**.
2. Enter a new value to override the inherited value.

## Viewing Inherited Values

When you configure DHCP properties that contain inherited values, the appliance displays the information based on the inheritance sources. The following table summarizes what the appliance can display:

When you see...	it means...	For details, see...
<b>Inherited From &lt;object&gt;</b>	the DHCP property has a definite value from an inheritance source.	Simple Inheritance below
<b>Inherited From Upper Level</b>	the appliance cannot determine the inherited value or inheritance source for the DHCP property.	Unknown Inheritance below
<b>Inherited From Multiple</b>	the DHCP property has the same value that it inherits from multiple sources.	Multiple Inheritance below
<b>Settings Inherited from Multiple Ancestors, View Multiple Inheritance Scenarios</b>	the DHCP property has multiple values that it inherits from multiple sources, and you can view the values and their corresponding sources by clicking the <b>View Multiple Inheritance Scenarios</b> link.	Multiple Inheritance below

## Simple Inheritance

When a DHCP property has an inherited value from a specific source, the appliance displays the value. It also displays **Inherited From <object>** (where <object> can be the Grid, member, network, shared network, or DHCP range) to indicate the source from which the value is inherited.

For example, when you set DHCP properties at the Grid level and do not override the properties at any level, the members, networks, shared networks, DHCP ranges, fixed addresses, reservations, host addresses, and roaming hosts inherit these properties from the Grid. The appliance displays the property value and **Inherited From Grid Infoblox** for each configured DHCP property, as shown in the figure Simple Inheritance below.

## Simple Inheritance

The screenshot shows the DHCP configuration interface for 'infoblox.localdomain (Member DHCP Properties)'. The 'Basic' tab is active. Under 'IPV4 PROPERTIES', the 'Authoritative' checkbox is unchecked, with the text 'DHCP server is authoritative' and 'Inherited from Grid Infoblox' below it. The 'Lease Time' is set to '12 Hours'. The 'Unlimited Lease Time' checkbox is unchecked, with a yellow warning box stating: 'Inadvertently selecting the Unlimited Lease Time check box or using this option incorrectly could cause a serious network outage in the future when all available leases are allocated'. Below this, the 'Microsoft Clients Code Page' is set to 'None', also with the text 'Inherited from Grid Infoblox'.

## Unknown Inheritance

In some cases, DHCP properties may not have definite inherited values and inheritance sources. The following are examples of unknown inheritance:

- The appliance cannot determine the inheritance sources of the DHCP properties in a template until you use the template to create an object.
- When a network or a DHCP range does not have an assigned member, it does not have a clear definition of an inheritance source because a network or a DHCP range inherits properties from a member.
- When individual networks in a shared network do not have member assignments, the shared network has unknown inheritance because the shared network inherits DHCP properties from a member and its networks.
- All roaming hosts have unknown inheritance because the DHCP properties can be inherited from different DHCP ranges within a network view.

In cases where the source of the inheritance is unknown, the appliance displays **Inherited From Upper Level** as the inheritance source. As shown in the figure Unknown Inheritance, network 10.1.1.0 has unknown lease time value because it does not have any assigned member.

## Unknown Inheritance

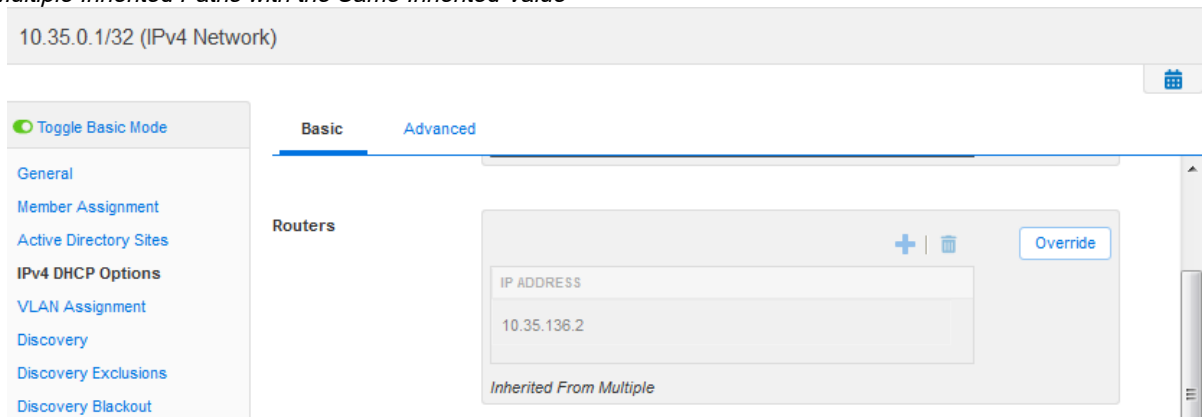
The screenshot shows the DHCP configuration interface for '10.35.0.1/32 (IPv4 Network)'. The 'Basic' tab is active. Under 'Lease Time', the value is empty, and the unit is set to 'Seconds'. The 'Unlimited Lease Time' checkbox is unchecked, with a yellow warning box stating: 'Inadvertently selecting the Unlimited Lease Time check box or using this option incorrectly could cause a serious network outage in the future when all available leases are allocated'. Below this, the text 'Inherited From Upper Level' is displayed.

## Multiple Inheritance

As illustrated in the figure *Multiple Inherited Paths with the Same Inherited Value*, a network can have multiple inherited values and inheritance sources for DHCP properties when it is served by multiple members. When an object inherits a DHCP property from different sources, the property value can be the same from all sources or it can be different. When the value is the same, the appliance displays the value in the property field. When there are multiple values inherited from multiple paths, the appliance displays the information to indicate so.

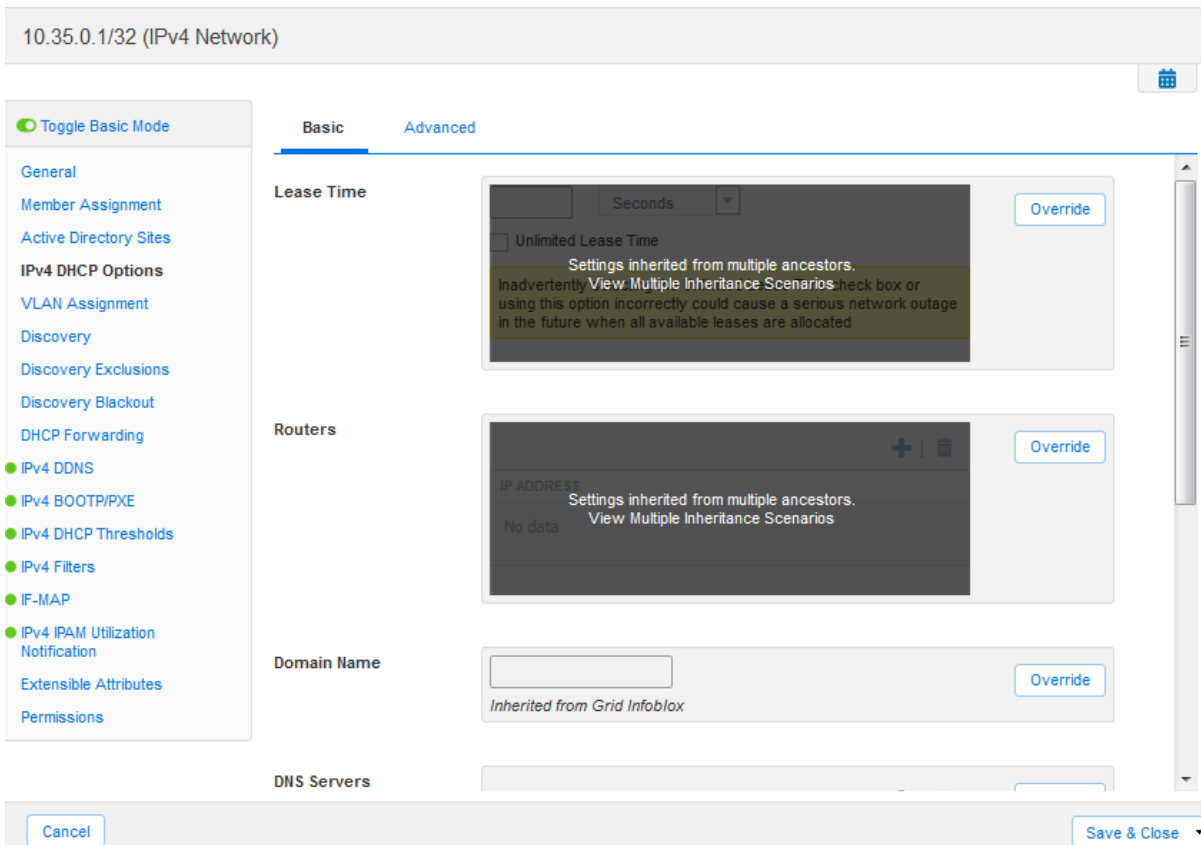
In a Grid, when two members serve the same network, the network inherits DHCP properties from both associated members. If both members have the same configured DHCP property, the network inherits the same value from both members. For example, when DHCP network 10.1.1.0 has two associated members and both members have the lease time set for 20 hours, the appliance displays the lease time value and **Inherited From Multiple** to indicate the value is inherited from multiple sources, as shown in the following figure.

### *Multiple Inherited Paths with the Same Inherited Value*



In the same Grid with the two members serving the same network, the network inherits different values for the same properties if you override the Grid configuration on one member but not on the other. For example, you can configure different PXE lease times for the members and configure a member as an authoritative DHCP server for the domain and the other not. In this case, the appliance displays **Settings inherited from multiple ancestors** and provides a **View Multiple Inheritance Scenarios** link so you can view the inherited values and paths, as shown in *Multiple Inheritance Sources with Multiple Values*.

### *Multiple Inheritance Sources with Multiple Values*



For example, to view the multiple inherited values of the **Authoritative** field, click **View Multiple Inheritance Scenarios**, and the *Multiple Inheritance Viewer* displays the inherited values from the two members. Since member1.foo.net does not have a configured value for this field, the viewer displays **Not Set**, as shown in the figure *Multiple Inheritance Viewer* below. You can use this information to determine whether you want to keep the inherited values or configure new ones.

#### Multiple Inheritance Viewer

Multiple Inheritance Viewer		
	INFOBLOX.LOCALDOMAIN	CLOUDMEMBER.COM
Lease Time	2147483647	43200

Another scenario of multiple inherited levels is when you have multiple DHCP properties that can inherit the same or multiple values from different sources. For example, when you configure multiple DHCP custom options, each of the options can inherit the same or multiple values from multiple paths. You can override the inherited options and configure new ones at a specific level other than the Grid level. Though these options are grouped under *DHCP Custom Options*, the appliance treats each of them as a separate property. The appliance groups the inherited options at the top, as shown in *DHCP Custom Options with Multiple Inheritance Sources*. You can override these options but you cannot delete them. For multiple values inherited from multiple sources, you can view the values in the *Multiple Inheritance Viewer* by clicking **View Inheritance**, as shown in the figure *Multiple Inheritance Viewers for Options*.

### DHCP Custom Options with Multiple Inheritance Sources

**Broadcast Address**  Override  
*Inherited from Grid Infoblox*

**Custom DHCP Options**

subnet-mask (1)  Override  
*Inherited From Multiple*

host-name (12) Multiple Members Override  
*Inherited From Multiple* [View Inheritance](#)

root-path (17)  Override  
*Inherited from Grid Infoblox*

Choose option  - +

### Multiple Inheritance Viewers for Options

Multiple Inheritance Viewer		
	INFOBLOX.LOCALDOMAIN	CLOUDMEMBER.COM
host-name (12)	hostname	hostname1

When you configure email notification for the Grid or Grid member from the **Data Management** tab -> **Grid** tab, the email address you enter there is inherited by the DHCP configuration for the Grid, members, networks, and DHCP ranges unless you override it at a specific level. The appliance uses this email address to send notification for a DHCP range when the DHCP usage crosses either the effective watermark threshold. For information, see [Configuring Thresholds for DHCP Ranges](#).

A network container inherits DHCP options from its parent and grandparent network containers. A network container does not inherit DHCP options defined at the Grid or member level.

Note the following about the DHCP option inheritance:

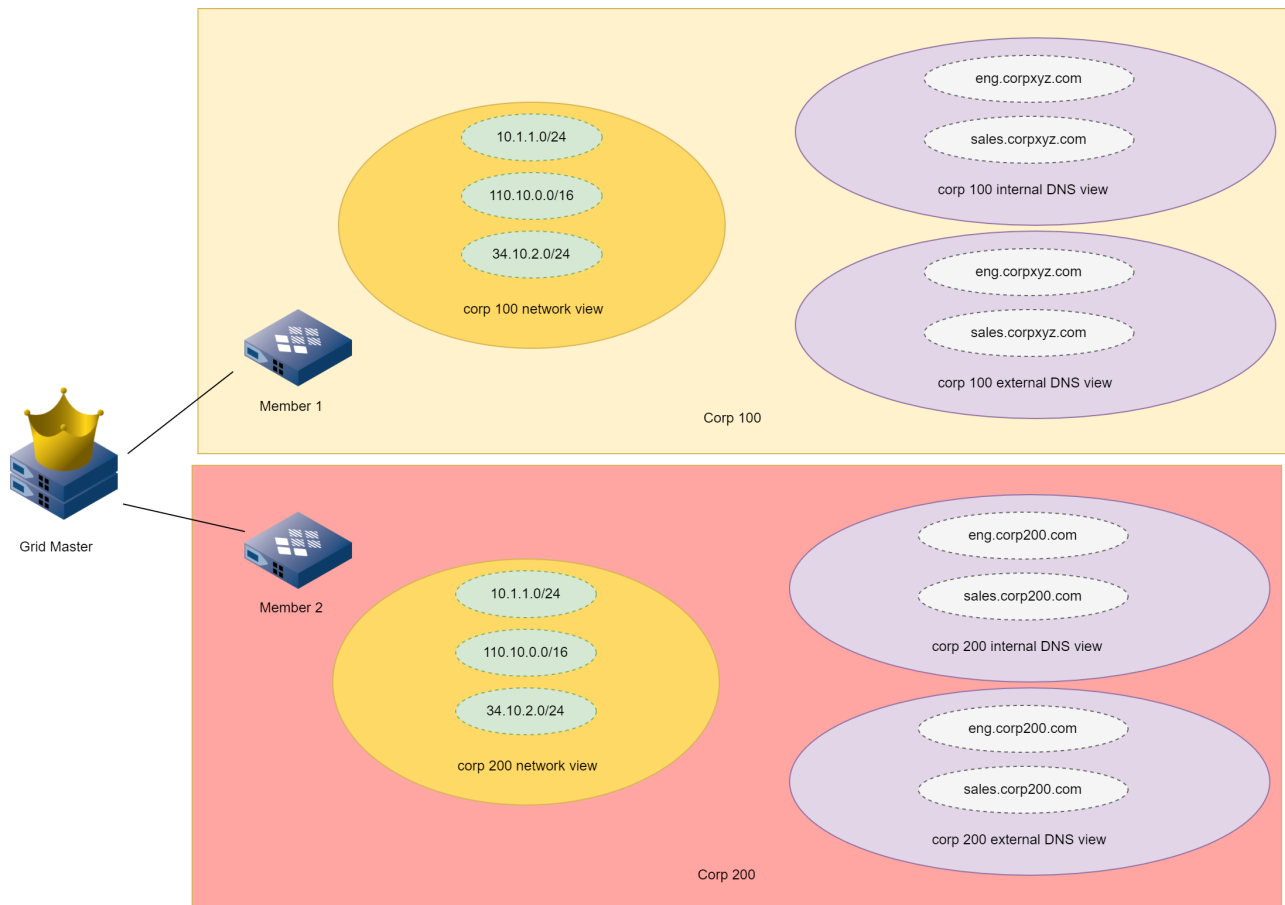
- For networks and shared networks, you can override an inherited DHCP option defined at the Grid or Member level.
- A shared network without a parent network container continues to inherit DHCP options from its parent Grid or member. The parent object is derived from the first network within the shared network.
- A network inherits DHCP option from its parent object. For example, if a network has a parent network container parent and parent shared network parent, if a DHCP option is overridden on the shared network, then this overridden value gets inherited. If the DHCP option is overridden on a network container, then this overridden value gets inherited. Otherwise, the network continues to inherit from its parent Grid or member.

## Configuring Network Views

A network view is a single routing domain with its own networks and shared networks. A network view can contain both IPv4 and IPv6 networks. All networks must belong to a network view.

You can manage the networks in one network view independently of the other network views. Changes in one network view are not reflected in other network views. Because network views are mutually exclusive, the networks in each view can have overlapping address spaces with multiple duplicate IP addresses without impacting network integrity. For example, two corporations, Corp 100 and Corp 200, merge. They each have their own networks and DNS domains. They also have their own private IP address spaces in the 10.0.0.0/24 network. Both corporations have DHCP and DNS servers, and use dynamic DNS updates. The DHCP servers of each corporation serve IP addresses for networks in their respective corporations. The DHCP clients in each corporation update DNS zones within their DNS domains. They plan to migrate the networks and hosts in Corp 200 to the Corp 100 address space and the corpxyz.com domain. To support both networks in the meantime and to facilitate the migration, you can configure an Infoblox Grid to centrally manage the networks and domains of both corporations. As shown in the figure *Two Network Views Managed by a Grid*, you can configure network views for each corporation and manage their networks independently of the other. Member 1 serves DNS and DHCP to Corp 100. The networks of Corp 100 are contained in the corp 100 network view, which is associated with both the internal and external DNS views of the corpxyz.com domain. Member 2 serves DNS and DHCP to Corp 200. The networks of Corp 200 are in the corp 200 network view, which is associated with both the internal and external DNS views of the corp200.com domain. The two corporations have one overlapping network, 10.1.1.0/24.

*Two Network Views Managed by a Grid*



A Grid member can serve one network view only, but a network view can be served by multiple Grid members. DHCP failover associations must be defined within a single network view, and both the primary and secondary peer must serve



the same network view.

The NIOS appliance provides one default network view. You can rename the default view and change its settings, but you cannot delete it. There must always be at least one network view in the appliance. If you do not need to manage overlapping IP address spaces in your organization, you can use the system-defined network view for all your networks. You do not need to create additional network views. But if there are overlapping IP address spaces and you need more than one network view, you can create up to 1000 network views.

Each network view must be associated with at least one DNS view. The default network view is always associated with the default DNS view, which also cannot be deleted. When you create a network view, the appliance automatically creates a corresponding DNS view with the same name as the network view, but with "default" prepended to the name. You can then rename that system-defined DNS view, but you cannot delete it.

A network view can be associated with multiple DNS views (as shown in the figure *Two Network Views Managed by a Grid*), but a DNS view cannot be associated with more than one network view. Each network view must be associated with a unique set of DNS views.

You can initiate a network discovery in only one network view at a time. When you run a discovery task, the appliance sends updates to all DNS views associated with the network view. (For information about network discoveries, see [IP Discovery and vDiscovery](#).)

## Adding Network Views

All networks must belong to a network view. You can use the default network view on the appliance and create additional network views, as needed. If you plan to enable DDNS (dynamic DNS) updates on any of the networks, DHCP ranges and fixed addresses in the network view, you must set parameters that specify which DNS view is updated for each network view.



### Note

If there are more than 20 network views, the appliance lists the available network views in the *Network View Selector* dialog box. If there are 20 or less than 20 network views, the appliance displays them in the drop-down list.

To create a network view:

1. From the **Administration** tab, select the **Network Views** tab, and then click the Add icon.
2. In the *Network View* wizard, do the following:
  - **Name:** Enter the name of the network view.
  - **Comment:** Enter useful information about the network view. The **Cloud** section displays if the Cloud Network Automation license is installed on the Grid Master. For information, see [Deploying Cloud Network Automation](#).
  - **Delegate authority from the Grid Master**
    - **Delegate To:** This field indicates whether the authority for the network view you want to create has already been delegated to a Cloud Platform Appliance. Click **Select** to choose the Cloud Platform Appliance to which you want to delegate authority. The *Member Selector* displays only Cloud Platform Appliances in the Grid. Click the member, and Grid Manager displays the member name next to this field. This cloud member now assumes authority for this network view, and the Grid Master does not have authority anymore. You can also click **Clear** to remove authority delegation from the selected Cloud Platform Appliance and return authority back to the Grid Master. Note that you need to install CP license before you configure a Grid with CP member. If you attempt to add a network view without installing CP license, the **Delegate To** field shows the following warning message:  
*There are no objects to select on the wizard.*  
NIOS does not display any warning message that prompts you to install the CP license.
  - **Restricting synchronization of network views**
    - **Disable sync to MGM:** Select this checkbox to disable synchronization. This will restrict the synchronization of all objects that are associated with the network view. This checkbox is available only on the managed Grid when it remains joined with the Multi-Grid Master.

3. Click **Next** to enter values for required extensible attributes or add optional extensible attributes for the network view. For information, see [About Extensible Attributes](#).
4. Click **Next**, and then save the configuration or select:

**Configure DDNS Properties:** Configure the DNS zones that are associated with the network view to receive DDNS updates. When you select this option, the *Configure DDNS Properties* dialog box appears. The appliance saves the network view entry before it opens the *Configure DDNS Properties* dialog box. For information, see [Configuring DDNS Updates](#).

## Modifying Network Views

1. From the **Administration** tab, select the **Network Views** tab -> *network\_view* checkbox, and then click the Edit icon.
2. The *Network View* editor provides the following tabs from which you can edit data:
  - **General:** You can modify **Name** and **Comments** in this tab. When the Cloud Network Automation license is installed on the Grid Master, Grid Manager displays the following in the **Cloud** section: You see the following when Cloud Network Automation is deployed (For information, see [Deploying Cloud Network Automation](#)):
    - **Cloud Usage:** This field indicates whether this object is associated with any specific cloud extensible attributes or within a scope of delegation. It can be one of the following:
      - **Cloud from adapter:** Indicates that this object has been created by a cloud adapter and it may or may not be within a scope of delegation at the moment.
      - **Cloud from delegation:** Indicates that this object is within the scope of delegation or the object itself defines a scope of authority delegation, and it is not created by a cloud adapter.
      - **Used by cloud:** Indicates that this network or network container is associated with the extensible attribute **Is External** or **Is Shared** and the value is set to True, which implies the network is a private or shared network managed by the CMP, and it is not **Cloud from adapter** or **Cloud from delegation**.
      - **Non-cloud:** The object is a regular NIOS object and is not within the scope of any authority delegation nor is it associated with any of these extensible attributes: **Cloud API Owned**, **Is External**, or **Is Shared**. NIOS admin users can modify this object based on their permissions.
    - **Owned By:** A cloud object can be owned by the Grid Master or the cloud adapter. When the object is created by the Cloud Platform member, this shows **Grid**. If the object is created by the cloud adapter, this shows **Adapter**.
    - **Delegated To:** This tells you whether a cloud object has been delegated to a Cloud Platform Appliance or not. If the cloud object has a parent object and the parent has been delegated, this field shows the parent delegation and you cannot modify the field.
  - **Restricting synchronization of network views**
  - **Disable sync to MGM:** Select this checkbox to disable synchronization. This will restrict the synchronization of all objects that are associated with the selected network view. This checkbox is available only on the managed Grid when it remains joined with the Multi-Grid Master.
- **Members:** This tab displays the members that provide DHCP services for the networks in this network view. You cannot modify information in this tab. It displays the following:
  - **Name:** The name of the DHCP member.
  - **IP Address:** The IP address of the DHCP member.
  - **Failover Association:** The name of the failover association to which the DHCP member belongs. If there are multiple failover associations, only the first one is displayed.
  - **Comment:** The information that you entered for the DHCP member. You can sort the information in the table by column. You can also print and export the information.
- **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network view. You can also modify the values of extensible attributes. For information, see [About Extensible Attributes](#).
- **Permissions:** This tab displays only if you belong to a superuser admin group. For information, see [Administrative Permissions for DHCP Resources](#).

## Deleting Network Views

You can delete any network view, except for the default network view. You can delete a network view that has only one DNS view associated with it. You cannot delete a network view that has more than one DNS view associated with it. When you delete a network view, the appliance deletes all the networks and records within the network view.

To delete a network view:

1. From the **Administration** tab, select the **Network Views** tab -> *network\_view* checkbox, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.

The appliance removes the network view and its associated DNS views. You can restore the network view from the Recycle Bin, if enabled. If you restore a network view, the appliance restores the associated DNS views as well. For information about the Recycle Bin, see [Using the Recycle Bin](#).



### Note

You cannot delete a network view that has a VLAN object assigned to it. For more information, see [Assigning VLANs to a Network](#).

## Configuring DHCP Properties

When you configure a NIOS appliance to function as a DHCP server, you can set DHCP properties that control how the appliance operates and enable DHCP service for IPv4 and IPv6.

You can also specify configuration information the appliance includes in its IPv4 and IPv6 DHCP messages. When a DHCP server assigns an IP address to a client, it can include information the client needs to connect to the network and communicate with other hosts and devices on the network. You can set these properties at the Grid level and override them for a member, network, shared network, DHCP range, fixed address, IPv4 reservation, host address, or roaming host.

When you configure a DHCP object that has inherited DHCP properties, you can either keep the inherited properties or override them. The appliance displays the inherited values and the levels from which the DHCP properties are inherited. For information, see [DHCP Inheritance](#).

This section explains how to configure DHCP IPv4 and IPv6 properties. It contains the following topics:

- [Configuring General IPv4 DHCP Properties](#)
- [Configuring Ping Settings](#)
- [Configuring DHCP Lease Management](#)
- [Ignoring DHCP Client Identifiers](#)
- [Configuring IPv4 BOOTP and PXE Properties](#)
- [About IPv4 DHCP Options](#)
- [Configuring Thresholds for DHCP Ranges](#)
- [Configuring DHCPv6 Properties](#)
- [About DHCPv6 Options](#)
- [Configuring DHCP IPv4 and IPv6 Common Properties](#)
- [Configuring DHCP Logging](#)
- [Configuring IF-MAP](#)
- [Starting DHCP Services on a Member](#)
- [Viewing DHCP Member Status](#)

## Configuring General IPv4 DHCP Properties

When you configure general IPv4 DHCP properties at the Grid level, the configuration applies to the entire Grid. Though you can set DHCP properties at the Grid level, you can enable DHCP services at the member level only. Infoblox recommends that you configure the DHCP properties before you enable DHCP on the appliance. Depending on the properties, you can override some of them for the members, networks, DHCP ranges, fixed addresses, reservations, host addresses, and roaming hosts. To override an inherited DHCP property, click **Override** next to the property to enable the configuration.

### Specifying Authoritative

Only authoritative DHCP servers can send clients DHCPNAK messages when they request invalid IP addresses. For example, a client moves to a new subnet and broadcasts a DHCPREQUEST message for its old IP address. An authoritative DHCP server responds with a DHCPNAK, causing the client to move to the INIT state and to send a DHCPDISCOVER message for a new IP address. Authoritative servers also respond to DHCPINFORM messages from clients that receive their IP addresses from the DHCP server and require additional options after the initial leases have been granted.

### Defining Lease Times

When you configure DHCP general properties, you can specify the length of time the DHCP server leases an IP address to a client. The default on the appliance is 12 hours, and you can change this default according to your network requirements. There are a number of factors to consider when setting the lease time for IP addresses, such as the types of resources and clients on the network, and impact to traffic and performance. With NIOS appliances, you can set lease times at different levels, based on these factors. You can also select the **Unlimited Lease Time** checkbox to grant unlimited lease time to IP addresses.



#### Warning

*Inadvertently selecting the Unlimited Lease Time checkbox could cause a network outage when the address range runs out of IP addresses.*

You can set a default lease time at the Grid level and then override this setting for specific members, network containers, networks, IP address ranges or fixed addresses when appropriate.

## Configuring General IPv4 Properties

To configure general IPv4 DHCP properties:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.  
**Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* checkbox, and then click the Edit icon.  
**Network Container:** From the **Data Management** tab, select the **IPAM** tab -> *network\_container* checkbox, and then click the Edit icon.  
**DHCP Range:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network->addr\_range* checkbox, and then click the Edit icon.  
**Fixed Address:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network->fixed\_address* checkbox, and then click the Edit icon.  
**Reservation:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network->reservation* checkbox, and then click the Edit icon.

2. In the *DHCP Properties* editor of a Grid or member, select the **General Basic** tab. For all other objects, click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, select the **IPv4 DHCP Options Advanced** tab.
3. Complete the following:
  - **Authoritative:** Select **DHCP server is authoritative** to set the DHCP server as authoritative for the domain. This can be set for the Grid, member, network container, network and range.
  - **Lease Time:** Enter the lease time and select the time unit from the drop-down list. The default is 12 hours.
    - **Unlimited Lease Time:** Select this option to set an infinite lease time for all IP addresses. To set all other properties for a Grid or member, toggle to the advanced mode and select the **General Advanced** tab to complete the following:
  - **Ignore Optionlist:** Select **Ignore optionlist requested by client and return all defined options** if you want the appliance to ignore the requested list of options in the DHCPREQUEST messages it receives from DHCP clients, and to include all the configured options in the DHCPACK and DHCP OFFER messages it sends back to the clients.
  - **LEASEQUERY:** Select **Allow LEASEQUERY** to enable the DHCP server to respond to DHCPLEASEQUERY messages.
4. Save the configuration and click **Restart** if it appears at the top of the screen.



#### Warning

*Inadvertently selecting the Unlimited Lease Time checkbox could cause a network outage when the address range runs out of IP addresses.*

## Configuring Fixed Addresses without Restarting the DHCP Service

When you configure or modify a fixed address, a DHCP service restart is required by default for the new configuration to take effect. You can override this default behavior by enabling the appliance to take immediate action when you configure or modify a fixed address outside a DHCP range without restarting DHCP service. This feature applies to host records for which DHCP is enabled, as it creates both the host address and the fixed address. You can enable this feature at the Grid or member level. Note that when you enable this feature, you cannot use the CLI command `set dhcp_expert_mode`.



#### Note

- Enabling this feature might have a significant performance impact on your appliance, depending on the number of fixed addresses you have configured.
- This feature works only for fixed addresses outside of a DHCP range. If you make a change to a fixed address inside a DHCP range, you must manually restart the DHCP service.

For Cloud Network Automation deployment, this feature is automatically enabled on the Cloud Platform Appliance that has a valid Cloud Platform license installed. For information about Cloud Network Automation, see [Deploying Cloud Network Automation](#).

To enable immediate fixed address configuration:

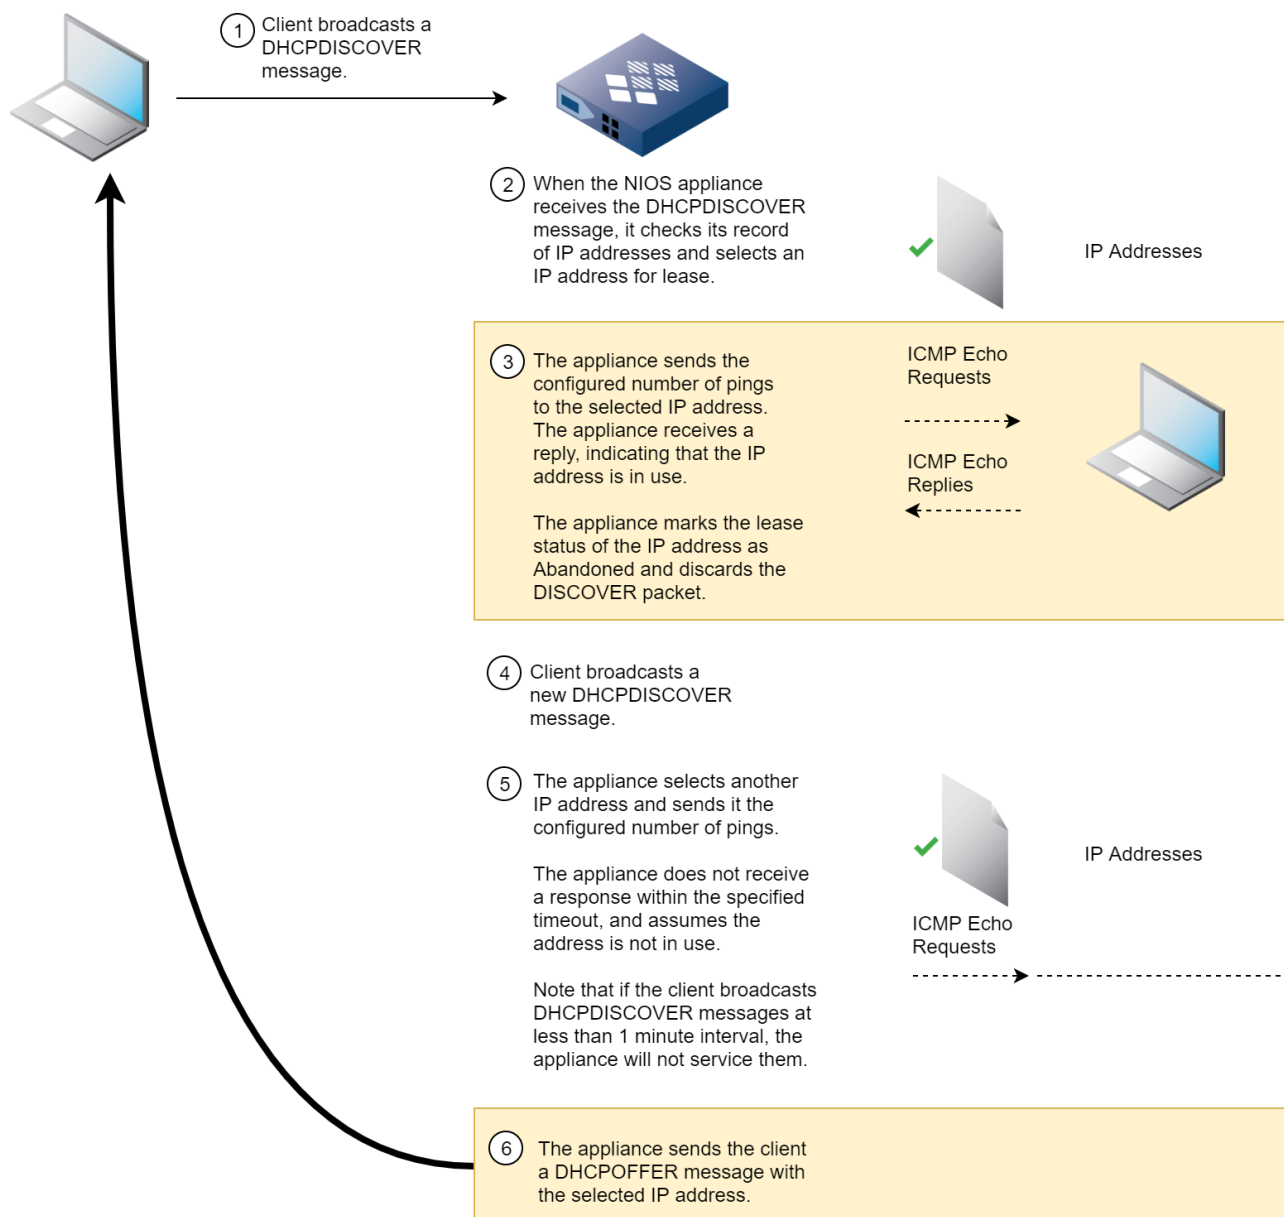
1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.
2. In the *DHCP Properties* editor, click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, click the **General** tab -> **Advanced** tab and complete the following:
  - **Immediate FA Configuration:** Select this checkbox to enable the new configuration immediately without restarting DHCP service when you modify or delete a fixed address.
3. Save the configuration and restart DHCP service.

## Configuring Ping Settings

When a DHCP client first tries to connect to a network, it broadcasts its request for an IP address. When the appliance receives such a request, it checks its record of assigned IP addresses and leases. Because there are a limited number of IP addresses available, the appliance reassigns IP addresses whose leases might have expired. Therefore, once the appliance selects a candidate IP address for lease, it sends an ICMP echo request (or ping) to the IP address to verify that it is not in use.

If the appliance receives a response, this indicates that the IP address is still in use. Note that the lease status for this IP address is **Abandoned**. Only after the DHCP client broadcasts a new request for an IP address that the appliance selects another candidate IP address and sends it a ping. This flow continues until the appliance finds an IP address that does not respond to the ping as depicted in the figure Ping Overview below. The appliance then sends a DHCP OFFER message with the unused IP address to the DHCP client.

Figure 26.1 Ping Overview





By default, the appliance pings the candidate IP address once and waits one second for the response. You can change these default settings to better suit your environment. Though you can increase the ping or timeout value to accommodate delays caused by problems in the network, increasing any of these values increases the delay a client experiences when acquiring a lease. You can also disable the appliance from sending pings by changing the number of pings to 0.

You can define ping settings for an entire Grid, and when necessary, define different ping settings for a member. Settings at the member level override settings at the Grid level.

To configure ping settings:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.
2. In the *DHCP Properties* editor, click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, click the **General** tab -> **Advanced** tab and complete the following:
  - **Number of Ping Requests:** Enter the number of pings the appliance sends to an IP address to verify that it is not in use. The range is 0 to 10, inclusive. Enter 0 to disable DHCP pings.
  - **Ping Timeout:** Select the ping timeout value from the drop-down list.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring DHCP Lease Management

When setting up DHCP properties, you can configure how the appliance handles lease management. For example, when a DHCP client moves from one network to another, the DHCP lease granted by the DHCP server in the previous network remains associated with the network device until the lease expires. When this happens to multiple clients, the address range could run out of IP addresses; thus, preventing other clients from retrieving an IP address. To avoid this, you can configure the DHCP server to release a lease when a client moves out of the network. Note that this is valid only if the client has included the client ID while requesting a lease. For example, a DHCP IPv4 client has included the client ID while requesting a lease from network A. This client then moves from network A to B. The DHCP server immediately releases the lease for reuse by another client on network A, based on the client ID. If the client does not include the client ID in the request, the lease on network A remains active until the lease expires.

You can also enable one-lease-per-client to ensure that each DHCP IPv4 client receives only one lease at any given time. When you enable one-lease-per-client and a DHCP client sends a DHCPREQUEST for a particular lease, the appliance releases other leases that the client holds, on the interface that the client is currently using.

Enabling one-lease-per-client is useful when you want to control the number of leases on your subnets and ensure that each DHCP client receives only one lease at a time. Typically, you enable one-lease-per-client for a DHCP client that moves around a lot within different subnets and uses long leases.

Note that this feature supports only DHCP IPv4 clients. When you configure lease management at the Grid level, all members inherit the setting. You can override the Grid setting for each member.

### Note

One-lease-per-client enables a single lease per client on a per member basis, not on a Grid wide basis. Lease information is not replicated among members. Note that you must restart the DHCP service for the changes to take effect.

To configure DHCP lease management:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.  
**Standalone DHCP:** From the **Data Management** tab, select the **DHCP** tab, expand the Toolbar and click **System DHCP Properties**.
2. In the *DHCP Properties* editor, click **Toggle Advanced Mode** if the editor is in basic mode. Click the **General** tab -> **Advanced** tab and select the respective radio button under **Multiple Lease Management**:



- **Release leases for clients with Client IDs:** Select this option to release only those leases that have a client ID when the clients move from one network to another before the leases expire. This is selected by default. The DHCP server does not release leases that do not have a client ID.
  - **Retain leases for clients with Client IDs:** Select this option to retain all the leases the client holds. The DHCP server retains all the leases either with or without a client ID. The amount of time taken by the DHCP server to find a lease for a client might increase if you use this option when the pools are almost full. Choosing this option might also increase the amount of active leases in the pools when the client moves from one network to another.
  - **Allow only one lease per client:** Select this option to enable one-lease-per-client per Grid member. This is valid for leases both with and without client IDs.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Ignoring DHCP Client Identifiers

You can set the DHCP server to ignore the UID (unique client identifier) and MAC address (hardware address) of a DHCP client when it places a request to the DHCP server for a new lease. When you configure the appliance to ignore the MAC address of a DHCP client, you can specify the list of MAC addresses. You can enter up to 10 MAC addresses to be ignored. The appliance ignores all the MAC addresses if you do not specify any MAC addresses. If the **Ignore Hardware Address** option is enabled and a DHCP client makes a request without a client UID for a new lease, then the appliance drops this request. This option is disabled by default. When you enable "Ignore DHCP Client ID" and a DHCP client sends a DHCPREQUEST for a lease, the DHCP server identifies the DHCP client using the physical MAC address of the appliance while the UID is ignored. The DHCP server then allocates an IP address based on the MAC address of the DHCP client.

For example, when a DHCP client places a request for a new lease, the DHCP server identifies the DHCP client with the MAC address and allocates the same IP address that was previously allocated for that MAC address.

You can define this feature at the Grid level, which is inherited at the member, shared network, IPv4 network and range level. This feature is disabled by default.



### Note

This feature is applicable only to dynamic leases and does not have any effect on the static lease generated for fixed addresses or roaming hosts.

To ignore the client identifier of DHCP clients:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the **Toolbar**.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> and click the **Members** tab -> *member* checkbox -> *Edit* icon.  
**Standalone DHCP:** From the **Data Management** tab, select the **DHCP** tab, and then click **System DHCP Properties**.  
**Shared Network Editor:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Shared Networks** section -> *shared\_network* checkbox, and then click the *Edit* icon.  
**IPv4 Network Editor:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* checkbox, and then click the *Edit* icon.  
**IPv4 Range Editor:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> click on the network address. Select the *IP address range* checkbox, and then click the *Edit* icon.
2. In the **DHCP Properties** editor, select the **General** tab -> **Advanced** tab (or click **Toggle Advanced Mode**) and then complete the following:
  - **Accept Client Identifier and MAC Address:** Select this checkbox to instruct the DHCP server to recognize MAC address and client UID of a DHCP client when it requests for a new lease.
  - **Ignore Client Identifier:** By default, this checkbox is not selected at the Grid level. Select this checkbox to ignore the client identifier of a DHCP client while placing a request to the DHCP server for a new lease. The DHCP server will only identify the MAC address and ignores the client identifier. DHCP clients requesting leases with different client UIDs receive the same IP address based on the MAC address. The

initial default state is inherited from the Grid level. Click **Override** to modify the inherited setting. To inherit the Grid settings, click **Inherit** at the member, IPv4 network and range, or shared network level.

- **Ignore MAC Address:** By default, this checkbox is not selected at the Grid level. Select this checkbox to ignore MAC address of a DHCP client while placing a request to the DHCP server for a new lease. To override the value that has been inherited from the Grid, click **Override**. Click the Add icon, the appliance adds a row to the table. Click the row and enter the MAC address to be ignored. You can also select a checkbox and click the Delete icon to delete the MAC address. To inherit the Grid settings, click **Inherit** at the member, IPv4 network and range, or shared network level.

3. Save the configuration and click **Restart** at the top of the screen.

### Limitations of the Ignore Client ID Feature on DHCP Failover Associations

- You cannot assign a DHCP range that has the ignore DHCP client ID feature enabled to a DHCP failover association if:
  - one of the members is an external DHCP server in the failover association.
  - one of the members is running a NIOS version earlier than 6.12.
- The DHCP failover association does not work if a DHCP range having multiple inherited values has the ignore DHCP client ID feature enabled on one server and disabled on the other.
- The range assigned to a DHCP failover association and the member (failover peer) must have the same DHCP range setting. The DHCP failover association does not work if a range associated with it does not have the same ignore DHCP client ID setting as the member.

### Configuring IPv4 BOOTP and PXE Properties

You can configure the DHCP server to support IPv4 clients that use BOOTP (bootstrap protocol) or that include the TFTP server name option and boot file name option in their DHCPREQUEST messages. You can specify the name or IP address of the boot server and the name of the file the host needs to boot.

In addition, you can configure the DHCP server to support hosts that use PXE (Preboot Execution Environment) to boot remotely from a server. When such a host starts up, it first requests an IP address so it can connect to a server on the network and download the file it needs to boot. After it downloads the file, the host reboots and sends another IP address request. To better manage your IP resources, set a different lease time for PXE boot requests. You can configure the DHCP server to allocate an IP address with a shorter lease time to hosts that send PXE boot requests, so IP addresses are not leased longer than necessary.

#### Note

When you assign a failover association to serve DHCP ranges and networks, NIOS denies dynamic BOOTP clients by default, regardless of whether you select or deselect the **Deny BOOTP Requests** option from Grid Manager. However, if the DHCP ranges or networks are assigned to a single DHCP server (not a failover association), NIOS does not automatically deny dynamic BOOTP clients. In this case, you must manually select the **Deny BOOTP Requests** option through Grid Manager to ensure that NIOS denies BOOTP requests to avoid problems such as receiving two IP addresses for the same network device.

You can configure BOOTP and PXE properties at the Grid level and override them for members, IPv4 network containers, IPv4 networks, DHCP ranges, fixed addresses, and reservations, host addresses, and roaming hosts. You cannot configure BOOTP and PXE properties for IPv6 DHCP objects.

To configure or override BOOTP and PXE properties:

1. **Grid Level:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member Level:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.  
**Network Level:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* checkbox, and then click the Edit icon.  
**Network Container:** From the **Data Management** tab, select the **IPAM** tab -> *network\_container* checkbox, and then click the Edit icon.  
**DHCP Range Level:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** ->

*network*-> *addr\_range* checkbox, and then click the Edit icon.

**Fixed Address Level:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *fixed\_address* checkbox, and then click the Edit icon.

**Reservation:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *reservation* checkbox, and then click the Edit icon.

2. In the *DHCPProperties* editor, select the **BOOTP/PXE** tab and complete the following:

- **PXE Lease Time:** Click **Override** and select **Enable PXE Lease Time** if you want the DHCP server to use a different lease time for PXE clients. You can specify the duration of time it takes a host to connect to a boot server, such as a TFTP server, and download the file it needs to boot. For example, set a longer lease time if the client downloads an OS (operating system) or configuration file, or set a shorter lease time if the client downloads only configuration changes. Enter the lease time for the preboot execution environment for hosts to boot remotely from a server.
- **Deny BOOTP Requests:** Select this checkbox to disable the BOOTP settings and deny BOOTP boot requests. If you assign DHCP ranges or networks to a single DHCP server (not a failover association), NIOS does not automatically deny dynamic BOOTP clients. In this case, you must select this option to ensure that NIOS denies BOOTP requests to avoid problems such as receiving two IP addresses for the same network device.
- Complete the following in the **BOOTP Settings** section:
  - **Boot File:** Enter the name of the boot file the client must download.
  - **Next Server:** Enter the IP address or hostname of the boot file server where the boot file is stored. Complete this field if the hosts in your network send requests for the IP address of the boot server. If the TFTP server is the NIOS appliance that is also serving DHCP, enter the IP address of the appliance.
  - **Boot Server:** Enter the name of the server on which the boot file is stored. Clients can request for either the boot server name or IP address. Complete this field if the hosts in your network send requests for the boot server name. If the TFTP server is the appliance that is also serving DHCP, enter the name of the appliance

 **Note**

Enter values in both the **Next Server** and **Boot Server** fields if some hosts on your network require the boot server name and others require the boot server IP address.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

Note that a few characters need manual escaping when you configure a DHCP boot file name, in order to keep the *dhcpd.conf* file consistent. If you do not use appropriate escape characters, then it might lead to a non working boot file name. The following characters require manual escaping:

- '\t' - Tabulation character
- '\r' - Carriage return
- '\n' - New line
- '\b' - Bell
- '\xYY' - YY hex-number (a-f, 0-9)

For example, if you set the 'Boot File' to:

```
'\x86\topdir\subdir\file.img'
```

You might need to enter `\x` and `\t` as the manual escape characters:

```
'\\x86\\topdir\\subdir\\file.img'
```

You can also specify all `\` as the manual escape character:

```
'\\x86\\topdir\\subdir\\file.img'
```

The above commands result in the underlying *dhcpd.conf* file:

```
'\x5cx86\x5ctopdir\\subdir\\file.img'
```

or

```
'\x5cx86\x5ctopdir\x5csubdir\x5cfile.img'
```

## About IPv4 DHCP Options

DHCP options provide specific configuration and service information to DHCP clients. These options appear as variable-length fields at the end of the DHCP messages that DHCP servers and clients exchange. For example, DHCP option 3 is used to list the available routers in the network of the client and option 6 is used to list the available DNS servers.

An option space is a collection of options. ISC (Internet Systems Consortium) DHCP has five predefined option spaces: dhcp, agent, server, nwip, and fqdn. The NIOS appliance supports only the predefined DHCP option space, which contains the industry standard options as well as additional options you can configure as needed:

- **Predefined options:** These are option codes 1 to 125. They are allocated by the IANA and defined by IETF standards. The DHCP server knows these standard options, and they are predefined on the server. You cannot redefine these options or delete them from the DHCP option space.
- **Custom options:** These are option codes 126 to 254. They are not defined by IETF standards and are available for private use. You can use these option codes to provide configuration or service information that none of the predefined options provide.

You can also create option spaces to define new groups of options. For example, you can create additional option spaces to define vendor specific options, which are encapsulated in option 43. When a DHCP client requests vendor specific options, it makes a request using the vendor identifier set in option 60 and a list of requested vendor specific options (option 43). The DHCP server then responds with the list of replies for the various options encapsulated into option 43.

Note that custom options defined in the DHCP option space are included in the options section of the DHCP messages that DHCP servers and clients exchange. Custom options defined in a user-defined option space are always encapsulated in option 43 in DHCP messages.

You can apply options globally at the Grid level, or more specifically at the member, network, network container, range, host and roaming host levels.

A network inherits DHCP options from its parent object. You can override the inherited DHCP options configured at the Grid or Member level for the networks and shared networks. If a network has a parent network container and parent shared network and if you override the DHCP options on the shared network, then the network inherits the shared network values. If you override the DHCP options at the network container level, then the network inherits the network container values. Otherwise, the network continues to inherit DHCP options from its parent Grid or member. A shared network without a parent network container continues to inherit DHCP options from its parent Grid or member. The parent object is derived from the first network within the shared network. A network container inherits DHCP options from its parent and grandparent network containers. A network container does not inherit DHCP options defined at the Grid or member level.

To override an inherited value, click **Override** next to it and complete the appropriate fields. When you click **Override**, the appliance displays the value inherited from its parent object (if any). If you do not set any value at the higher level, the appliance displays the default DHCP options. The following table lists the default DHCP Options:

### Default DHCP Options

	Name	Default Value
IPv4 Common DHCP Options	Enable PXE Lease Time	Disabled
	Lease Time	43200
	Routers	Empty List

	Name	Default Value
	Domain Name	Empty
	DNS Servers	Empty List
	Broadcast Address	127.0.0.1
	Custom DHCP Options	""
	Lease Scavenging	Disabled
	Authoritative (Advanced)	Disabled
	Lease Deletion (Advanced)	Disabled
	Ignore Optionlist (Advanced)	Disabled
IPv4 DDNS	Enable DDNS Updates	Disabled
	DDNS Domain Name	""
	DDNS Update TTL	0
	DDNS Update Method	Interim
	Generate Hostname	Disabled
	Fixed Address Updates	Disabled
	Option 81 Support	Disabled
	Lease Renewal Update	Disabled
IPv4 Threshold Options	Enable DHCP Thresholds	Disabled
	High - Trigger	95
	High - Reset	85
	Low - Trigger	0
	Low - Reset	10
	Enable SNMP Warnings	Disabled

	Name	Default Value
	Enable Email Warnings	Disabled
	Email Addresses	Empty List
IPv4 BOOTP/PXE	Enable PXE Lease Time	Disabled
	Lease Time (Value)	0
	Deny-BOOTP-Requests	Disabled
	Boot File	""
	Next Server	""
	Boot Server	""
IPv6 DHCP Options	Valid Lifetime	43200
	Preferred Lifetime	27000
	Domain Name	""
	DNS Servers	
	Custom DHCP Options	
	Lease Deletion	Enabled
	Lease Scavenging	Disabled
IPv6 DDNS Options	Enable DDNS Updates	
	DDNS Domain Name	""
	DDNS Update TTL	0
	DDNS Update Method	Interim
	Generate Hostname	Disabled
	FQDN Support <ul style="list-style-type: none"> <li>• DHCP server always updates DNS</li> <li>• DHCP server updates DNS if requested by client</li> </ul>	Disabled

	Name	Default Value
	Lease Renewal Update	Disabled

You can also create an option filter the appliance uses to filter address requests by the DHCP options of requesting hosts. The filter instructs the appliance to either grant or deny an address request if the requesting host matches the filter. For information, see [Defining Option Filters](#).

The DHCP option configuration conforms to the following RFCs:

- *RFC 2132, DHCP Options and BOOTP Vendor Extension*
- *RFC3046, DHCP Relay Agent Information Option*. The supported options include option 60 (Client Identifier), 21 (Policy Filter), 22 (Maximum Datagram Reassembly Size), 23 (Default IP Time-to-Live), and 82 (Support for Routed Bridge Encapsulation).
- *RFC3925, Vendor-Identifying Vendor Options for Dynamic Host Configuration Protocol version 4 (DHCPv4)*
- *RFC2939, Procedures and IANA Guidelines for Definition of New DHCP Options and Message Types*

## DHCP Option Data Types

Each DHCP option is identified by a name and an option code number, and specifies a data type. The data type for some options is predefined. For example, in the DHCP option space, the data type for option 1: subnet-mask is an IP address. You cannot change the data type for this option. The data type for some options is user-defined and can be in one of the formats shown in the below table.

### DHCP Option Data Types

Data type	Specifies
String	An ASCII text string (the same as the text data type) or a list of hexadecimal characters separated by colons  Formatting to distinguish an ASCII text string from a hexadecimal string is important. For details, see the following section
Boolean	A flag with a value of either true or false (or on or off )
IP address	A single IP address
Array of IP addresses	A series of IP addresses, separated by commas  You can optionally include a space after each comma
Text	An ASCII text string
8-, 16-, or 32-bit unsigned integer	A numeric range of the following possible values  8-bit unsigned integer: from 0 to 255  16-bit unsigned integer: from 0 to 65,535  32-bit unsigned integer: from 0 to 4,294,967,295
8-, 16-, or 32-bit signed integer	A numeric range of the following possible values  8-bit signed integer: from -128 to 127  16-bit signed integer: from -32,768 to 32,767  32-bit signed integer: from -2,147,483,648 to 2,147,483,647



Data type	Specifies
Domain-list	A list of domain names, separated by spaces

When defining a hexadecimal string for a DHCP option (such as option 43, vendor encapsulated options), use only hexadecimal characters (0-9, a-f, or A-F) without spaces and separated by colons. The accepted form for a hexadecimal string, as presented in a regular expression, is `[0-9a-fA-F]{1,2}(:[0-9a-fA-F]{1,2})*`

Two examples of correctly written hexadecimal strings:

- aa:de:89:1b:34
- 1C:8:22:A3 (Note that the DHCP module treats a single hexadecimal character, such as "8" as "08".)

A few examples of incorrectly written hexadecimal strings:

- :bb:45:d2:1f – Problem: The string erroneously begins with a colon.
- bb:45:d2:1f: – Problem: The string erroneously ends with a colon.
- bb:4 5:d2:1f – Problem: The string erroneously includes a space between two characters ("4" and "5").
- bb:45:d2:1g – Problem: The string erroneously includes a nonhexadecimal character ("g").

The DHCP module treats incorrectly written hexadecimal strings as simple text strings, not hexadecimal strings. If the string appears in quotes, it is a text string.

## Configuring IPv4 DHCP Options

To use DHCP options, you can do the following:

- Define basic DHCP options, as described in the next section [Defining IPv4 DHCP Options](#).
- Configure one or more option spaces, as described in the next section [Defining IPv4 Option Spaces](#).
- Define custom options in the predefined DHCP option space or add options to an option space that you configured. For more information, see [Configuring Custom DHCP Options](#) below.
- Specify values for the options and apply them to the Grid, or to a member, network, range, fixed address, reservation, host, or roaming host. For more information, see [Applying DHCP Options](#) below.

## Defining IPv4 DHCP Options

You can define basic DHCP options that the DHCP server uses to provide configuration information to DHCP clients. The server includes these options in its DHCP messages.

To define DHCP options:

1. **Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* check box, and then click the Edit icon.

**Network Container:** From the **Data Management** tab, select the **IPAM** tab -> *network\_container* check box, and then click the Edit icon.

**DHCP Range:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network->DHCP\_range* check box, and then click the Edit icon.

**Fixed Address:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network->fixed\_address* check box, and then click the Edit icon.

**Reservation:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network->reservation* check box, and then click the Edit icon.

**Host Address:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network->host\_record* check box, and then click the Edit icon. Select the host IP address, and then click the Edit icon.

**Roaming Host:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Roaming Hosts** -> *roaming\_host* check box, and then click the Edit icon.

2. In the *DHCP Properties* editor, select the **IPv4 DHCP Options** tab and complete the following:
  - **Routers:** Click the Add icon. Grid Manager adds a row to the table. In the table, enter the IP address of the router that is connected to the same network as the DHCP client. When configuring this for a template,

enter the offset value of the IP address of the router. The DHCP server includes this information in its DHCPOFFER and DHCPACK messages.

- **DomainName:** Enter the name of the domain for which the Grid serves DHCP data. The DHCP server includes this domain name in Option 15 when it responds with a DHCPOFFER packet to a DHCPDISCOVER packet from a client. If DDNS is enabled on the DHCP server, it combines the host name from the client and this domain name to create the FQDN (fully-qualified domain name) that it uses to update DNS. For information about DDNS, see [Configuring DDNS Updates](#). When overriding the domain name already set by a parent object, enter the new value for the selected option or use "" to clear the value.
- **DNS Servers:** Click the Add icon. Grid Manager adds a row to the table. In the table, enter the IP address of the DNS server to which the DHCP client sends name resolution requests. The DHCP server includes this information in the DHCPOFFER and DHCPACK messages.
- **Broadcast Address:** Enter the broadcast IP address of the network to which the DHCP server is attached. When configuring this for a template, enter the offset value of the broadcast IP address of the network.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Defining IPv4 Option Spaces

DHCP members support the DHCP option space by default. You can create additional option spaces to provide additional configuration or service information. Note that custom options defined in a user-defined option space are always encapsulated in option 43 in DHCP messages

To add a custom option space:

1. From the **Data Management** tab, select the **DHCP** tab -> **Option Spaces** tab.
2. Click the Add icon -> **IPv4 Option Space**.
3. In the *Option Space* wizard, do the following:
  - **Name:** Enter the name of the option space.
  - **Comment:** Enter useful information about the option space.
  - **Options:** Click the Add icon to add options. For additional information, see the next section, [Configuring Custom DHCP Options](#).
4. Save the configuration and click **Restart** if it appears at the top of the screen.

After you create an option space and add options to it, you can apply the options as described in [Applying DHCP Options](#) below.

## Configuring Custom DHCP Options

You can define custom options in the DHCP option space or in an option space that you configured, as follows:

1. From the **Data Management** tab, select the **DHCP** tab -> **Option Spaces** tab.
2. Select either the **DHCP** option space or an IPv4 option space that you configured, and then click the Edit icon.
3. In the *Option Space* editor, click the Add icon to add a custom option. In the new row, complete the following:
  - **Name:** Enter the name of the custom DHCP option.
  - **Code:** Select an option code from the drop-down list. Select a number between 126 and 254 if you are adding custom options to the **DHCP** option space. If you are adding custom options to an IPv4 option space you configured, you can enter a number between 1 and 254.
  - **Type:** Select the option type (such as ip-address, text, boolean, and string as described in the DHCP Option Data Types table above).  
For example, to create an option that defines the IP addresses of Solaris root servers, enter the name SrootIP4, select option code 126, and then select the type as ip-address.  
Click the Add icon to add more options.
4. Save the configuration.

## Applying DHCP Options

Some options may apply to all networks and some may apply to specific ranges and even hosts. When you apply an option, you select the object to which the option is applied, such as the Grid member, or network, and then specify a

value for the option.

Use the following guidelines when specifying option values:

- Enter **false** or **true** for a Boolean Flag type value.
- Enter an ASCII text string, or enter a series of octets specified in hex, separated by colons.
- Separate multiple values by commas. For example, to enter multiple IP addresses for netbios-name-servers, enter a comma between each IP address.

Here are some examples of option names and correctly formatted values:

Option name	Value	Comment
option 61 dhcp-client-identifier	MyPC	Double quotes are no longer needed for string type values
dhcp-client-identifier	43:4c:49:45:54:2d:46:4f:4f	Series of octets specified in hex, separated by colons for a Data-string type value
netbios-name-servers	10.1.1.5,10.1.1.10	Multiple IP addresses separated by commas
option-80	ABC123	Custom option number 80 set to the string ABC123.

To apply DHCP options:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the Edit icon.  
**Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* check box, and then click the Edit icon.  
**Network Container:** From the **Data Management** tab, select the **IPAM** tab -> *network\_container* check box, and then click the Edit icon.  
**DHCP Range:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr\_range* check box, and then click the Edit icon.  
**Fixed Address:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *fixed\_address* check box, and then click the Edit icon.  
**Reservation:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *reservation* check box, and then click the Edit icon.  
**Host Address:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *host\_record* check box, and then click the Edit icon. Select the host IP address, and then click the Edit icon.  
**Roaming Host:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Roaming Hosts** -> *roaming\_host* check box, and then click the Edit icon.
2. In the **DHCP Properties** editor, select the **IPv4 DHCP Options** and complete the following:
  - The **Custom DHCP Options** section displays two fields. The first field displays **Choose option**. Click the arrow and select an option from the list. In the second field, enter a value for the selected option. Note that certain options have predefined data types and their values must be entered in a specific format. For information about the data types, see as described in the DHCP Option Data Types table above.
  - Click **+** to add another option, or click **-** to delete a previously specified option. When overriding an option already set by a parent object, enter the new value for the selected option or use "" to clear the value. Note that if you created an option space as described in Defining IPv4 Option Spaces above, this section displays a list of option spaces in the first drop-down menu, so you can select the option space of the option you want to define.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuration Example: Defining a Custom Option

In this example, you configure two custom options in the DHCP option space, and apply them to a DHCP range in the network 192.168.2.0/24.

Add the custom options to the DHCP options space:

1. From the **Data Management** tab, select the **DHCP** tab -> **Option Spaces** tab.
2. Select the **DHCP** check box, and click the Edit icon.
3. In the *DHCP (Option Space)* editor, click the Add icon. In the new row, complete the following:
  - **Name:** Enter **tftp-server**.
  - **Code:** Enter **150**.
  - **Type:** Select **array of ip-address**.
4. Click the Add icon to add another option. In the new row, complete the following:
  - **Name:** Enter **pxe-configfile**.
  - **Code:** Enter **209**.
  - **Type:** Select **text**.
5. Click **Save & Close**.

Enter values for the newly defined custom options and apply them to a DHCP range:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** subtab, and click the 192.168.2.0/24 network.
2. Click the 192.168.2.10 - 100 check box, and then click the Edit icon.
3. In the *DHCP Properties* editor, select the **DHCP** tab and complete the following in the **Custom DHCP Options** section:
  - From the drop-down list of options, select **tftp-server (150) array of address**. In the second field, enter **192.168.1.2**.  
Click **+** to add another option.
  - From the drop-down list of options, select **pxe-configfile (209) text**. In the second field, enter **pxe.config**, which is the file name of the boot image.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

The member then includes options 150 and 209 in its DHCP messages to clients that are allocated IP addresses from the DHCP range 192.168.2.10 - 100.

## Defining Option 60 Match Rules

The appliance uses option 60 (vendor-class-identifier) to forward client requests to the DHCP server for services that the clients require. You can define option 60 match rules and filter on these rules. You can set these rules for the Grid and override for a member.

To define option 60 for the Grid or member:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* check box, and then click the Edit icon.
2. In the *DHCP Properties* editor, click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, click the **DHCP Options** tab -> **Advanced** tab and complete the following:  
To override the Grid configuration for a member, click **Override** next to the property. Grid Manager hides the Grid configuration. You can then add new values for the member.
  - a. **Option60 (Vendor Class Identifier) Match Rules:** Click the Add icon if you want to add a match rule to a vendor class option. The appliance adds a row to the table. Complete the following:
    - **Option Space:** Select an option space from the drop-down list. This field appears only when you have custom option spaces. The appliance uses the default **DHCP** option space if you do not have custom option spaces.
    - **Match Value:** Enter the value you want the appliance to use when matching vendor class options.
    - **Is Substring:** Select this check box if the match value is a substring of the option data.
    - **Substring Offset:** Enter the number of characters at which the match value substring starts in the option data. Enter 0 to start at the beginning of the option data, enter 1 for the second position,

and so on. For example, when you enter 2 here and have a match value of RAS, the appliance matches the value RAS starting at the third character of the option data.

- **Substring length:** Enter the length of the match value. For example, if the match value is SUNW, the length is 4.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

## About the DHCP Relay Agent Option (Option 82)

A typical relationship between a DHCP client, relay agent, and the NIOS appliance on a network is as follows:

1. A DHCP client broadcasts a DHCPDISCOVER message on its network segment.
2. A DHCP relay agent on that segment receives the message and forwards it as a unicast message to one or more DHCP servers (such as NIOS appliances).
3. If the NIOS appliance accepts the address request, it responds to the relay agent with a DHCPOFFER message. If the appliance denies the request, it does not send any response in case other DHCP servers that might be involved respond instead.
4. The relay agent forwards the response to the client, usually as a broadcast message.

The situation is different for individual hosts connecting to the Internet through an ISP, usually over a circuit-switched data network.

1. A host connects to its ISP's circuit access concentration point, authenticates itself, and requests an IP address.
2. The circuit access unit relays the address request to a DHCP server, which responds with a DHCPOFFER message.

To avoid broadcasting the DHCPOFFER over the network segment on which the host made the request, the relay agent sends the response directly to the host over the established circuit.

Option 82 assists the agent in forwarding address assignments across the proper circuit. When a relay agent receives a DHCPDISCOVER message, it can add one or two agent IDs in the DHCP option 82 suboption fields to the message. The relay agent IDs are:

- **Circuit ID:** This identifies the circuit between the remote host and the relay agent. For example, the identifier can be the ingress interface number of the circuit access unit (perhaps concatenated with the unit ID number and slot number). The circuit ID can also be an ATM virtual circuit ID or cable data virtual circuit ID.
- **Remote ID:** This identifies the remote host. The ID can be the caller ID telephone number for a dial-up connection, a user name for logging in to the ISP, a modem ID, and so on. Because the remote ID is defined on the relay agent, which is presumed to have a trusted relationship with the DHCP server, and not on the untrusted DHCP client, the remote ID is also presumably a trusted identifier.

### Note

For information about the relay agent option, refer to *RFC3046, DHCP Relay Agent Information Option*.

In addition to the relay agent IDs, NIOS also supports the Option 82 Link Selection and Server ID Override sub-options, which allow DHCPv4 to operate in a network architecture where direct communication between the DHCP server and DHCP client is undesirable or infeasible. You can configure these sub-options to direct DHCP traffic to go through the relay agent and have more control over your DHCP communications.

The Link Selection sub-option provides a mechanism to separate the subnet/link in which the DHCP client resides from the GIADDR (Gateway IP address). The GIADDR field in a DHCP message is populated by the relay agent and is typically used to inform the DHCP server about the subnet in which the DHCP client resides and to inform the DHCP server of the IP address to use to communicate with the relay agent. In situations where the GIADDR might not be the appropriate subnet from which IP addresses should be allocated, you can use the Link Selection sub-option to explicitly set the subnet from which IP addresses are allocated to the client.

The Server ID Override sub-option allows the relay agent to tell the DHCP server what IP address, instead of the server's address, must be used in the response. Generally, the response from the server contains the IP address of the DHCP server itself in the Server-ID option. You can use the Server ID Override sub-option to specify a new value for the server ID that is inserted in the reply packet by the DHCP server. Configuring the Server ID Override sub-option allows the relay agent to have the clients send all unicast messages to the relay agent instead of the DHCP server.

 **Note**

If you want the Link Selection and Server ID Override sub-options to be included in the DHCP relayed messages, you must configure them on the DHCP relay agent. You cannot configure them on NIOS. For more information about these sub-options, refer to <https://tools.ietf.org/html/rfc3527> and <https://tools.ietf.org/html/rfc5107>.

On the NIOS appliance, you can do the following with option 82:

- Screen address requests through a relay agent filter you set up using option 82. For more information, see [About Relay Agent Filters](#).
- Use the relay agent information (circuit ID or remote ID) as a host identifier when configuring a fixed address, though you cannot do so in a host record. For information about how to configure a circuit ID or remote ID as an identifier, see as described in [Adding IPv4 Fixed Addresses](#).
- Define how Grid Manager displays the relay agent ID, circuit ID, and remote ID (when applicable) in the detailed lease information panel. For information about how to configure the logging format for option 82, see [Defining Logging Format for DHCP Option 82](#) below.

### Defining Logging Format for DHCP Option 82

When you define the circuit ID or remote ID of the relay agent as a host identifier, you can choose the logging format Grid Manager uses to display the IDs in the detailed lease information panel. For information about viewing lease information, see [Viewing Detailed Lease Information](#).

To define logging format for the agent ID, circuit ID and remote ID, complete the following:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the **Toolbar**.  
**Independent Appliance:** From the **Data Management** tab, select the **DHCP** tab, and then click **System DHCP Properties**.
2. In the *Grid DHCP Properties* or *System DHCP Properties* editor, select the **General** tab -> **Advanced** tab.
3. Select one of the following for **Logging format for Option 82:**
  - **Hexadecimal:** When you select this, Grid Manager displays the agent ID, circuit ID, and remote ID in hexadecimal format in the detailed lease information. This is the default format.
  - **Plaintext:** When you select this, Grid Manager displays the agent ID, circuit ID, and remote ID in plain text in the detailed lease information.

 **Note**

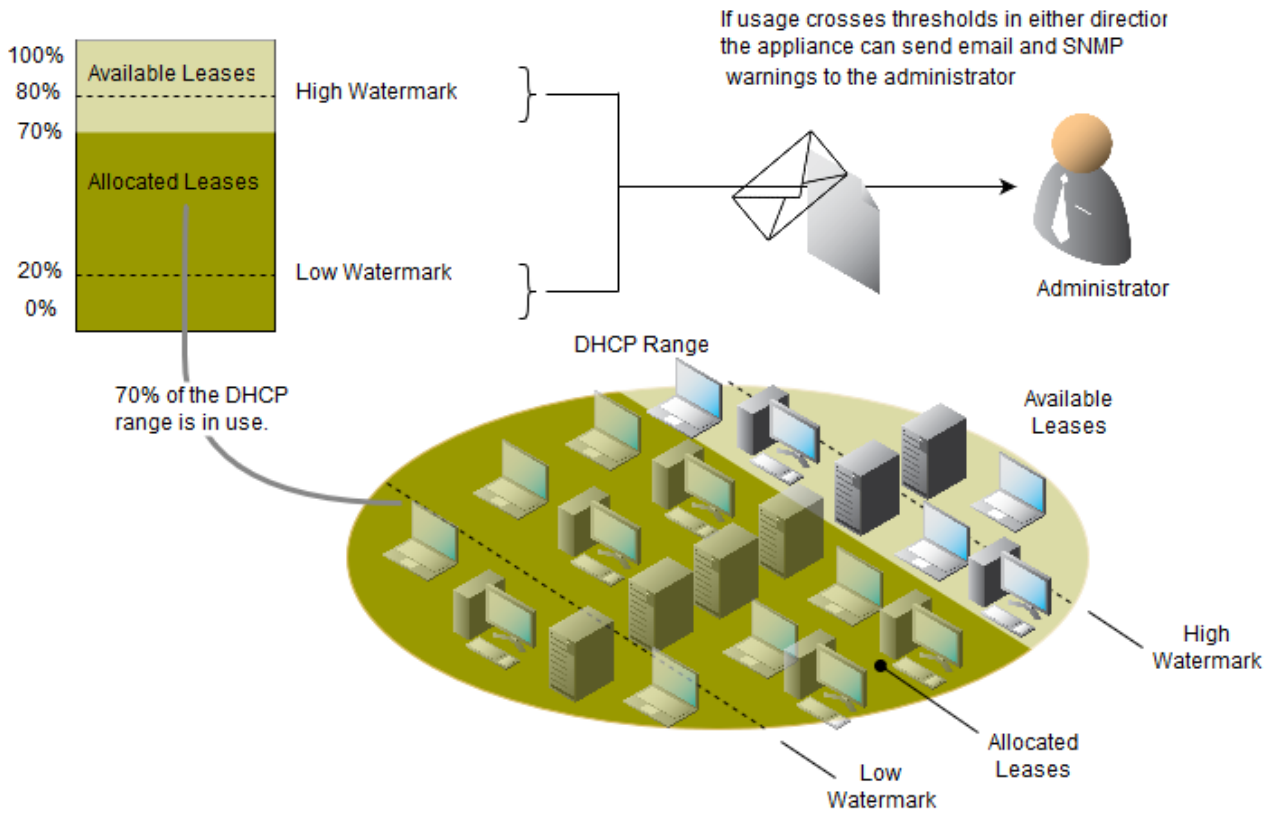
You cannot override this Grid setting at the member level. Also, changing the logging format requires a DHCP service restart.

## Configuring Thresholds for DHCP Ranges

Grid Manager can provide a view of the current overall DHCP range usage for the DHCP ranges defined on each Grid member. The view is in the form of a percent: address leases in use/total addresses for each network. Such information can indicate if there is a sufficient number of available addresses at each of these levels. It can also provide information about the distribution of address resources, indicating if there are too many unused addresses in one location while all the addresses in another are in use.

In addition to viewing the percent of addresses in use, you can also apply high and low thresholds for each DHCP range. These watermarks represent thresholds above or below which DHCP range usage is unexpected and might warrant your attention. For example, usage falling below a low threshold might indicate network issues preventing the renewal of leases. When usage for a DHCP range crosses a threshold, the appliance makes a syslog entry and — if configured to do so — sends the administrator alerts as SNMP traps and email notifications. The figure Overall DHCP Address Usage for a DHCP Range below illustrates the relationship of allocated and available addresses to high and low watermarks in a DHCP range.


*Overall DHCP Address Usage for a DHCP Range*



You can define watermarks at the Grid, member, network, and DHCP range levels, but the appliance applies them solely to DHCP ranges. Because the appliance applies settings hierarchically in a parent-child structure, by defining watermarks once at a higher level, DHCP ranges can then inherit these settings without your needing to redefine them for each range. For example, if you set high and low watermarks for a Grid, then each Grid member, each network, and each DHCP range inherits these settings. However, if you override these settings at the member level, then the network and DHCP ranges for that member inherit its settings. If you override the Grid member settings at the network level, then that network and any DHCP ranges within that network inherit the network-level settings. Finally, you can set high and low watermarks for an individual DHCP range, which override anything set at a higher level. The below figure shows different high and low watermark settings at different levels. Although you can set thresholds at four levels (Grid, Grid member, network, and DHCP range), the NIOS appliance applies them to DHCP ranges.

*High and Low Watermarks*





Address usage in a DHCP range can trigger an event and an email notification when it crosses a watermark. You must enable DHCP threshold and email warnings to receive events and notifications. The following are actions that do and do not trigger an address usage event and notification:

Address usage triggers an event and the appliance sends a notification when the percentage of the allocated addresses in the DHCP range:

- Exceeds the high watermark
- Drops below or equals to the high watermark after exceeding it
- Drops below the low watermark
- Exceeds the low watermark after dropping below it

Address usage does not trigger an event when the percentage of the allocated addresses in the DHCP range:

- Never exceeds the low watermark
- Initially exceeds the low watermark
- Reaches a watermark but does not cross it

**Note**

You can effectively disable address usage events for a DHCP range by setting its high watermark at 100% and the low watermark at 0% (default setting for the low watermark). Because address usage cannot cross these watermarks, no events can occur.

You can configure the threshold settings at the Grid level and override them at the member, network, and DHCP range levels. To override an inherited DHCP property, click **Override** next to the property to enable the configuration. For information, see [Overriding DHCP Properties](#).

To configure thresholds:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member*

checkbox, and then click the Edit icon.

**Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* checkbox, and then click the Edit icon.

**Network Container:** From the **Data Management** tab, select the **IPAM** tab -> *network\_container* checkbox, and then click the Edit icon.

**DHCP Range:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr\_range* checkbox, and then click the Edit icon.

- In the editor, select the **IPv4 DHCP Thresholds** tab and complete the following:
  - DHCP Thresholds:** Specify the following:
    - Enable DHCP Thresholds:** Select **Enable DHCP Thresholds** to enable the DHCP threshold feature.
      - High:** Enter a number between 0 and 100. Enter Trigger and Reset values. If the percentage of allocated addresses in a DHCP range exceeds the Trigger value, the appliance makes a syslog entry and—if configured to do so—sends an SNMP trap and an email notification to a designated destination. When the percentage first reaches the Reset value after it hit the Trigger value, the appliance sends an SNMP trap and an email notification to a designated destination. The default Trigger value is 95, and the default Reset value is 85.
      - Low:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range drops below the Trigger value, the appliance makes a syslog entry and—if configured to do so—sends an SNMP trap and an email notification to a designated destination. When the percentage first reaches the Reset value after it hit the Trigger value, the appliance sends an SNMP trap. The default Trigger value is 0 and the default Reset value is 10.
    - Enable SNMP Warnings:** Select this for the appliance to send an SNMP trap to the trap receiver that you define for the Grid when the address usage in a DHCP range crosses a high or low mark threshold.
    - Enable Email Warnings:** Select this for the appliance to send an email notification to an administrator if the address usage in a DHCP range crosses a high or low mark threshold.
  - Email Addresses:** Click **Override** to override the Grid administrator email address configured in the **Data Management** tab -> **Grid** tab. This address is not hierarchically inherited from the Grid DHCP configuration. Click the Add icon, and then enter an email address to which you want the appliance to send email notifications when the DHCP range for the network crosses a threshold. You can create a list of email addresses.
- Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring DHCPv6 Properties

The following sections describe how to configure properties and options that apply to DHCPv6 objects only. You can configure and define the following DHCP properties:

- General properties, as described in the next section, see [Defining General IPv6 Properties](#).
- DHCP options, as described in [About DHCPv6 Options](#).

### Defining General IPv6 Properties

You can configure general DHCPv6 properties at the Grid level and override them at the member and lower levels.

- Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
  - Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.
  - Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* checkbox, and then click the Edit icon.
  - Network Container:** From the **Data Management** tab, select the **IPAM** tab -> *network\_container* checkbox, and then click the Edit icon.
  - Fixed Address:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network*-> *Fixed address* checkbox, and then click the Edit icon.

2. In the *DHCP Properties* editor, select the **IPv6 DHCP Options** tab, and complete the following:
  - **Valid Lifetime:** Specify the length of time addresses that are assigned to DHCP clients remain in the valid state. When this time expires, an address becomes invalid and can be assigned to another interface.
  - **Preferred Lifetime:** Specify the length of time that a valid address is preferred. A preferred address can be used with no restrictions. When this time expires, the address becomes deprecated.
  - **Domain Name:** Enter the name of the domain for which the Grid serves DHCP data.
  - **DNS Servers:** Click the Add icon. Grid Manager adds a row to the table. In the table, enter the IPv6 addresses of DNS recursive name servers to which the DHCP client can send name resolution requests. The DHCP server includes this information in the DNS Recursive Name Server option in Advertise, Rebind, Information-Request, and Reply messages.
  - **Custom IPv6 DHCP Options:** In the first field, select one of the following from the drop-down list:
    - **DHCPv6:** Select this to apply DHCPv6 options.
    - **DHCP:** Select this to apply DHCP options (dhcp-renewal-time or dhcp-rebinding-time).
  - Click the **Choose option** drop down and then select an option from the list.
  - In the third field, enter a value for the selected option. Note that certain options have predefined data types and their values must be entered in a specific format. Click + to add another option, or click - to delete a previously specified option. When overriding an option, enter the new value for the selected option.
3. Save the configuration.

## About DHCPv6 Options

DHCPv6 options provide configuration and service information to IPv6 clients. Just like IPv4 options, IPv6 options appear as variable length fields at the end of the DHCPv6 messages.

Just as in IPv4, the NIOS appliance supports the following options in the DHCPv6 options space:

- **Predefined options:** These are the option codes defined in RFC 3315. You cannot redefine these options or delete them from the DHCP option space. Option codes 1-48 are reserved and cannot be used to define custom options.
- **Custom options:** These are option codes 49 to 254. They are not defined by IETF standards and are available for private use. You can use these option codes to provide configuration or service information that none of the predefined options provide.

You can also create option spaces to define new groups of options. For example, you can create additional option spaces to define vendor specific options, which are encapsulated in DHCPv6 option 17. When an IPv6 client requests vendor specific options, it makes a request using the vendor specific options (option 17). The DHCP server then responds with the list of replies for the various options encapsulated into option 17.

Note that custom options defined in the DHCP option space are included in the options section of the DHCP messages that DHCP servers and clients exchange.

You can apply options globally at the Grid level, or more specifically at the member, network, range, host and roaming host levels.

## Configuring DHCPv6 Options

To use DHCPv6 options, you can do the following:

- Configure one or more option spaces, as described in the next section [Defining IPv6 Option Spaces](#).
- Define custom options in the predefined DHCPv6 option space or add options to an option space that you configured. For more information, see [Configuring Custom DHCP Options](#) below.
- Specify values for the options and apply them to the Grid, or to a member, network, fixed address, host, or roaming host. For more information, see [Applying DHCPv6 Options](#) below.

## Defining IPv6 Option Spaces

DHCP members support the DHCPv6 option space by default. You can create additional option spaces to provide additional configuration or service information.

To add a custom option space:

1. From the **Data Management** tab, select the **DHCP** tab -> **Option Spaces** tab.

2. Click the Add icon -> **IPv6 Option Space**.
3. In the *IPv6 Option Space* wizard, do the following:
  - **Name:** Enter the name of the option space.
  - **Enterprise Number:** Enter the vendor's Enterprise Number that is registered with IANA.
  - **Comment:** Enter useful information about the option space.
  - **Options:** Click the Add icon to add options. For additional information, see the next section, *Configuring Custom DHCP Option*
4. Save the configuration and click **Restart** if it appears at the top of the screen.

After you create an option space and add options to it, you can apply the options as described in [Applying DHCP Options](#).

## Configuring Custom IPv6 DHCP Options

You can define custom options in the DHCP option space or in an option space that you configured, as follows:

1. From the **Data Management** tab, select the **DHCP** tab -> **Option Spaces** tab.
2. Select either the **DHCPv6** option space or an IPv6 option space that you configured, and then click the Edit icon.
3. In the *Option Space* editor, click the Add icon to add a custom option. In the new row, complete the following:
  - **Name:** Enter the name of the custom DHCP option.
  - **Code:** Enter a number from 1 to 65535 to add a custom option in the DHCP option space or in an IPv6 option space that you have configured.
  - **Type:** Select the option type (such as ipv6-address, text, boolean, and string as described in [DHCP Option Data Types](#) table).

Click the Add icon to add more options.

4. Save the configuration.

## Applying DHCPv6 Options

You can apply some options at the Grid or member level, and some options to specific networks, shared networks, fixed addresses and roaming hosts. When you apply an option, you select the object to which the option is applied, such as the Grid, member, or network, and then specify a value for the option.

Use the following guidelines when specifying option values:

- Enter **false** or **true** for a Boolean Flag type value.
- Enter an ASCII text string, or enter a series of octets specified in hex, separated by colons.
- Separate multiple values by commas. For example, to enter multiple IP addresses for netbios-name-servers, enter a comma between each IP address.

DHCPv6 options support the same data types as DHCP IPv4 options. For more information about the data types, see [DHCP Option Data Types](#). To apply DHCP options:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then select **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.  
**Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* checkbox, and then click the Edit icon.  
**Network Container:** From the **Data Management** tab, select the **IPAM** tab -> *network\_container* checkbox, and then click the Edit icon.  
**Fixed Address:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *fixed\_address* checkbox, and then click the Edit icon.  
**Host Address:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *host\_record* checkbox, and then click the Edit icon. Select the host IP address, and then click the Edit icon.  
**Roaming Host:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Roaming Hosts** -> *roaming\_host* checkbox, and then click the Edit icon.
2. In the *DHCP Properties* editor, select the **IPv6 DHCP Options** and complete the following:
  - **Custom IPv6 DHCP Options:** In the first field, select one of the following from the drop-down list:

- **DHCPv6**: Select this to apply DHCPv6 options.
- **DHCP**: Select this to apply DHCP options (dhcp-renewal-time or dhcp-rebinding-time).

In the second field, click the **Choose option** arrow and select an option from the list. In the third field, enter a value for the selected option. Note that certain options have predefined data types and their values must be entered in a specific format. For information about the data types, see [DHCP Option Data Types](#).

Click **+** to add another option, or click **-** to delete a previously specified option. When overriding an option, enter the new value for the selected option.

Note that if you created an option space, this section displays a list of option spaces in the first drop-down menu, so you can select the option space of the option you want to define.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Limitations of Custom DHCP Options

Note the following limitations when you enable DHCPv6 rapid commit:

- If you enable DHCPv6 rapid commit at a higher hierarchy level, all child objects at the lower hierarchy levels inherit this option. You cannot override this option for any child objects.
- You cannot add a network with rapid commit enabled to a shared network or define rapid commit in a network that belongs to a shared network.
- If the client does not include the DHCPv6 Rapid Commit option in the SOLICIT messages, the server will respond with an ADVERTISE using the default four-message exchange instead of REPLY using the two-message exchange even though rapid commit is enabled for the DHCP range in which the address is allocated.

## Configuring DHCP IPv4 and IPv6 Common Properties

This section describes DHCP properties that apply to both IPv4 and IPv6. It includes the following sections:

- [Configuring UTF-8 Encoding for Hostnames](#)
- [Associating Networks with Zones](#)
- [Keeping Leases in Deleted IPv4 and IPv6 Networks and Ranges](#)
- [Configuring Fixed Address Leases For Display](#)
- [Scavenging Leases](#)
- [DHCPv6 Lease Affinity](#)

### Configuring UTF-8 Encoding for Hostnames

When you configure the appliance as a DHCP server, the appliance supports UTF-8 encoding of hostnames that are encoded with Microsoft Windows code pages. You can configure the DHCP services on the appliance to convert these client hostnames to UTF-8 characters. The appliance stores the UTF-8 encoded hostnames in the database. If you also configure the DHCP services on the appliance to perform DDNS updates, the appliance sends the UTF-8 encoded host names in the DDNS updates. You can configure the UTF-8 encoding of host names at the Grid DHCP service and member DHCP service levels. For information on UTF-8 encoding, see [Printing from Grid Manager](#).

The appliance displays the host names in their original characters in the following:

- DHCP lease history
- DHCP lease details
- IP address management
- Syslog
- Audit log

To configure UTF-8 encoding for hostnames:

1. **Grid**: From the **Data Management** tab, select the **DHCP** tab, and then select **Grid DHCP Properties** from the Toolbar.  
**Member**: From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.
2. In the *DHCP Properties* editor, select the **General Basic** tab and complete the following:

- **IPv4 Properties**
    - **Microsoft Clients Code Page:** From the drop-down list, select the code page with which the host names are encoded when the appliance converts the Microsoft code page encoded host names to UTF-8 characters.
  - **IPv6 Properties**
    - **Microsoft Clients Code Page:** From the drop-down list, select the code page with which the host names are encoded when the appliance converts the Microsoft code page encoded host names to UTF-8 characters.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Associating Networks with Zones

You can associate IPv4 and IPv6 networks with DNS zones to limit the zones that admins can use when they create DNS records for IP addresses in the networks. When a network is associated with one or more zones and an admin creates a DNS record for one of its IP addresses, Grid Manager allows the admin to create the DNS record in the associated zones only. For example, if you associate the 10.1.0.0/16 network with the corpxyz.com zone, admins are allowed to create DNS records in the corpxyz.com zone only for IP addresses in the 10.1.0.0/16 network; or if you associate the 2001:db8:1::/48 network with the corp200.com zone, admins are allowed to create DNS records in the corp200.com zone only for IP addresses in the 2001:db8:1::/48 network.

This feature applies to A, AAAA and host records only. It does not apply to records in a shared record group. If you are creating a host record with multiple IP addresses in different networks, the networks must be associated with the zone of the host record.

If a network is not associated with a zone, admins can create DNS records for its IP addresses only in zones with no network associations as well.

You can associate a network with any authoritative zone whose primary server is a Grid member or a Microsoft server, or is unassigned. You cannot associate networks with zones that have external primary servers.

You can associate a network with multiple zones, and associate a zone with more than one network. You can associate IPv4 and IPv6 network containers and networks with zones. When you associate a network container with zones, its networks inherit the zone associations. You can override the zone associations at the network level.

If you split a network, the resulting subnets inherit the zone associations. If you join networks, the resulting network retains the zone associations of the network that you selected when you performed the join operation. You can override the inherited zone associations of individual networks. Subzones do not inherit the network associations of their parent zones.

When you import data into a zone that is associated with a list of networks, the imported A, AAAA and host records must have IP addresses in the associated networks. Grid Manager does not allow you to import A, AAAA and host records with IP addresses in unassociated networks.

When you associate a network with a zone, the DNS records created before the association are not affected. But if you edit an A, AAAA or host record after the association, Grid Manager does not allow you to save the record if its IP address is not in an associated network.

To associate an IPv4 or IPv6 network with a zone:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* checkbox, and then click the Edit icon.
2. In the *DHCP Network* editor, click **Toggle Advanced Mode** if the editor is in basic mode.
3. When the additional tabs appear, click the **Advanced** subtab of the **General** tab.
4. Click the Add icon and select the zone you want to associate with the network.
  - Optionally, select a default zone. When you create or edit an A, AAAA or host record from a network in the **IPAM** tab, Grid Manager automatically selects the default zone that is assigned to the network.
5. Save the configuration or click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Viewing the Networks Associated with a Zone

You can view the IPv4 or IPv6 networks associated with a zone from the zone editor. The tab to display network associations in zone editors is visible only if the primary server is a Grid member, a Microsoft server, or unassigned.

To view the network associations of a zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox, and then click the Edit icon.
2. In the *Authoritative Zone* editor, click **Toggle Advanced Mode** if the editor is in basic mode.

3. When the additional tabs appear, click the **Advanced** subtab of the **General** tab.

The Network Associations table lists the networks and their corresponding comments. You cannot change the network associations in this editor. Navigate to the *DHCP Network* editor of the network, to change the zone associations.

### Keeping Leases in Deleted IPv4 and IPv6 Networks and Ranges

You can configure the DHCP server to store leases in a deleted DHCP range for up to one week after the leases expire. When you add a new DHCP range that includes the IP addresses of these leases or assign the DHCP range to another member within the Grid, the appliance automatically restores the active leases. You can configure this feature for the Grid, and override the configuration for members, networks, and DHCP ranges.

To keep active leases in a deleted DHCP range:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.  
**Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* checkbox, and then click the Edit icon.  
**Network Container:** From the **Data Management** tab, select the **IPAM** tab -> *network\_container* checkbox, and then click the Edit icon.  
**Range:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *range* checkbox, and then click the Edit icon.
2. In the *DHCP Properties* editor of the Grid or member, click **Toggle Advanced Mode** if the editor is in basic mode, and then click the **General** tab -> **Advanced** tab. In the Network editor or Range editor, click **Toggle Advanced Mode** if the editor is in basic mode, and then click **IPv4DHCPOptions** -> **Advanced** or **IPv6 DHCP Options** -> **Advanced**. Complete the following:
  - **IPv4 Properties**
    - **Lease Deletion:** When you select **Keep leases from deleted range until one week after expiration** and delete a DHCP range with active leases, the appliance stores these leases for up to one week after they expire.
  - **IPv6 Properties**
    - **Lease Deletion:** When you select **Keep leases from deleted range until one week after expiration** and delete a DHCP range with active leases, the appliance stores these leases for up to one week after they expire.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

### Configuring Fixed Address Leases For Display

You can configure the DHCP server to capture the hostname and lease time of a fixed address when you assign an IPv4 or IPv6 fixed address to a client. The appliance displays the hostname, and the start and end time of each fixed address lease in the *Current Leases* panel in Grid Manager.

You can set this at the Grid level only for IPv4 and IPv6 leases.



1. From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, click the **General** tab -> **Advanced** tab and complete the following:
  - **IPv4 Properties**
    - **Fixed Address Lease:** Select **Capture hostname and lease time when assigning Fixed Addresses**. The appliance displays the host name, and the start and end time of each fixed address lease in the *Current Leases* panel. If there are multiple records (A, host, and lease) for the IP address, it also displays the information for the records. This option is available in the Grid Properties editor only.)
  - **IPv6 Properties**
    - **Fixed Address Lease:** Select **Capture hostname and lease time when assigning Fixed Addresses**. The appliance displays the host name, and the start and end time of each fixed address lease in the *Current Leases* panel. If there are multiple records (AAAA, host, and lease) for the IP address, it also displays the information for the records. This option is available in the Grid Properties editor only.)
3. Save the configuration.

## Scavenging Leases

The accumulation of free and backup DHCPv4 leases; and free, expired, and released DHCPv6 leases results in unnecessary growth of database objects. The DHCP lease scavenging feature enables member DHCP servers to automatically delete free and backup IPv4 leases; and free, expired, and released IPv6 leases that remain in the database beyond the specified period of time, thus reducing the number of database objects.

When you enable this feature for DHCPv4 leases, the appliance permanently deletes the free and backup IPv4 leases, and you can no longer view or retrieve the lease information. This option can be enabled globally at the Grid level, and more specifically for a member, shared network, network, network container, DHCP range, network template, DHCP range template.

When you enable this feature for DHCPv6 leases, the appliance permanently deletes the free, expired, and released IPv6 leases, and you can no longer view or retrieve the lease information. This option can be enabled at the Grid level, and overridden at the member level.

The period of time that you specify is the duration after the expiration date of a lease, not its release date. For example, you specify a time period of 5 days when you enable this feature. If the lease time of an IP address is 10 days, but the lease is released after five days, the appliance still deletes the lease from the database after 15 days because the IP address has been leased.

Note that If you plan to enable this feature after upgrading from a previous NIOS version, Infoblox recommends that you enable it during off-peak hours, as it may impact DHCP services.

To enable scavenging of IPv4 and IPv6 leases:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.  
**Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* checkbox, and then click the Edit icon. This is applicable for IPv4 lease scavenging only.  
**Network Container:** From the **Data Management** tab, select the **IPAM** tab -> *network\_container* checkbox, and then click the Edit icon. This is applicable for IPv4 lease scavenging only.  
**DHCP Range:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr\_range* checkbox, and then click the Edit icon. This is applicable for IPv4 lease scavenging only.
2. In the editor, click **Toggle Advanced Mode** if the editor is in basic mode, and then click the **General** tab -> **Advanced** tab.

In the Network editor for IPv4 lease scavenging, click **Toggle Advanced Mode** if the editor is in basic mode, and then click **IPv4 DHCP Options** -> **Advanced**.

Complete the following:

- **IPv4 Properties**
  - **Lease Scavenging:** This is disabled by default. Select the **Scavenge free and backup leases after** checkbox and specify the number of days or weeks that free and backup IPv4 leases remain in

the database before they are automatically deleted. This can be set for the Grid, member, network, and network container.

- **IPv6 Properties**

- **Lease Scavenging:** This is disabled by default. Select the **Scavenge free, expired and released leases after** checkbox and specify the number of hours, days, or weeks that free, expired, and released IPv6 leases remain in the database before they are automatically deleted. The minimum is 6 hours and the maximum is 180 days. The default is one week. This can be set at the Grid and member level.

3. Save the configuration.

## DHCPv6 Lease Affinity

DHCPv6 ranges are usually large and the DHCPv6 server randomly selects a new lease each time a client requests for a lease. The client can use the lease until it expires. After its expiration, the lease stays in the database with an expired state. These expired leases eventually lead to the increase in the number of database objects, because the probability of expired IPv6 leases getting reused is low.

Infoblox provides a DHCPv6 lease affinity feature that allows you to reuse expired IPv6 leases for DHCP clients. When you enable this feature, the DHCPv6 server automatically renews the expired leases. A DHCP client can retrieve the same lease from the DHCPv6 server after it expires and retains the same IP address. This feature helps reduce the amount of IPv6 leases in the database as the DHCP server issues the same lease multiple times for the same client. The appliance ignores expired leases that are older than the specified period. Such leases are scavenged. Note that the grace period you define for lease scavenging is applicable for DHCPv6 lease affinity also. The minimum time period is six hours, maximum is 180 days and the default is set to seven days. For more information about scavenging leases, see the *Scavenging Leases* section.

The DHCPv6 server offers the same lease for a DHCP client, identified by DUID, after the lease expires and before the end of the grace period. The appliance removes the expired leases that are older than the grace period from the database.

DHCPv6 lease affinity and DHCPv6 lease scavenging are complementary features. For example, consider a scenario in which a visiting user gets an IPv6 lease that is retained for days, weeks, or months depending on the needs and then the user leaves. If the user returns and the lease is still within the grace period, the user gets the same IPv6 lease. This is lease affinity. When the user leaves, the IPv6 lease becomes inactive. This lease is scavenged after the grace period. Note the following about DHCPv6 lease affinity:

- It does not consider expired leases that are older than the grace period.
- It ignores expired leases that do not match known ranges.
- If no existing lease is found, then the DHCPv6 server finds a suitable expired lease that is not older than the grace period, which matches the client DUID and range configuration.
- The impact of the feature on the performance depends on the amount of expired DHCPv6 leases.
- When you activate the feature at the Grid level, it affects all underlying layers of inheritance.
- You cannot enable DHCPv6 lease affinity at the Grid and member levels during a scheduled full upgrade.
- DHCPv6 lease affinity remembers only permanent addresses and does not remember temporary addresses and prefix delegations.
- If the DHCPv6 range is out of available addresses when you enable DHCPv6 lease affinity, then the DHCP server tries to reuse the best abandoned lease, which indicates the lease that was abandoned longest time ago. If there are no such leases in the pool, the DHCP server reuses the best expired lease, which indicates the lease that expired longest time ago. This means that the expired lease becomes active and it is associated with the new client while the DHCP server removes any previous associations of the corresponding lease. Note that this happens only when the DHCPv6 range does not have any available addresses and there are no suitable abandoned leases.

To enable DHCPv6 lease affinity:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.

2. In the editor, click **Toggle Advanced Mode** if the editor is in basic mode, and then click the **General** tab -> **Advanced** tab.  
Complete the following:
  - **IPv6 Properties**
    - **Lease Scavenging**
      - **Remember client association for expired members:** This is disabled by default. Select the checkbox to remember and reuse expired IPv6 leases that are associated with DHCP clients. You can select this checkbox only when you select the **Scavenge free, expired and released leases after** checkbox. This can be set at the Grid and member levels.

Note that the appliance stores the leases, which are either deleted or removed, in the recycle bin. These leases then become free and are automatically dissociated from their clients. For example, if you delete a range accidentally and restore it again, the IPv6 leases associated with the respective range are no longer associated with the same set of clients.
3. Save the configuration.

## Configuring DHCP Logging

If you have a syslog server operating on your network, you can specify in which facility you want the server to display the DHCP logging messages. You can also select the Grid member on which you want to store the DHCP lease history log, as described in the next section *Configuring the Lease Logging Member*. You can configure DHCP and lease logging only on the Grid and member levels.

To specify DHCP logging for the Grid or member:

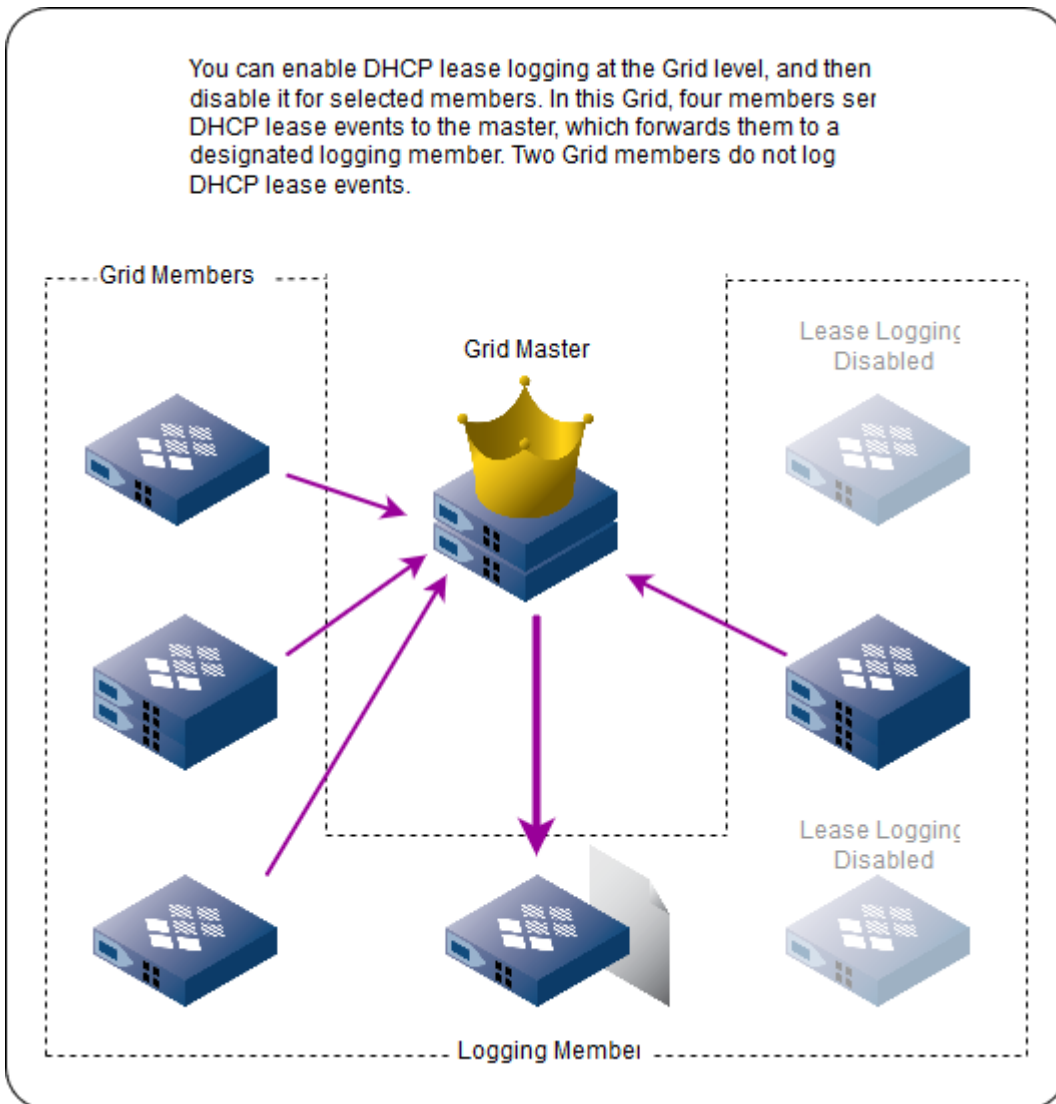
1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.
2. In the *DHCP Properties* editor, select the **Logging Basic** tab and complete the following:
  - **Syslog Facility:** From the drop-down list, select the facility that is used to tag syslog messages from the DHCP server. This facility can be used to filter messages on a central syslog server.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring the Lease Logging Member

Logging DHCP lease events makes significant CPU demands, especially when there is heavy DHCP activity. Therefore, Infoblox strongly recommends that you designate a Grid member other than the master as a logging member whenever possible. Another way to manage the increased load that logging introduces is to log selectively per Grid member. For example, you might want to log DHCP leases for members serving critical parts of your network and not keep historical logs for members serving other parts.

### *DHCP Lease History Logging with Member Overrides*

You can enable DHCP lease logging at the Grid level, and then disable it for selected members. In this Grid, four members send DHCP lease events to the master, which forwards them to a designated logging member. Two Grid members do not log DHCP lease events.



To specify lease logging for a member:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.
2. In the **Logging** tab, complete the following:
  - **Lease Logging:** Select **Enable Lease History** (for Grid) or **Log Lease Events from DHCP Server** (for member) to enable DHCP lease logging. To disable DHCP lease logging, clear the checkbox. You can set member overrides if you want to enable or disable lease logging per member.
  - **Send leases to:** For Grid only. Click **Select**. In the *Select Member* dialog box, select the Grid member on which you want to store the DHCP lease history log. Infoblox recommends that you dedicate a member other than the Grid Master as a logging member. If possible, use this member solely for storing the DHCP lease history log. If you do not select a member, no logging can occur. You can click **Clear** to remove the selected Grid member and select a new one.
3. Save the configuration and click **Restart** if it appears at the top of the screen.
4. For information about viewing current leases, see [Viewing Current Leases](#).

## Configuring IF-MAP

You can configure Infoblox DHCP servers to publish DHCP data to an IF-MAP server. The IF-MAP server takes real-time information from different sources and stores it in a shared database from which clients can retrieve information about network devices, their status and activities. For details about the IF-MAP protocol, refer to <http://www.trustedcomputinggroup.org>. For information about the Infoblox IF-MAP server, refer to the *Infoblox Administrator Guide for Infoblox Orchestration Server*.

Each Infoblox DHCP server in a Grid can function as an IF-MAP client, with the ability to publish lease information to an IF-MAP server. For information about how to configure an IF-MAP client, see [Configuring Members as IF-MAP Clients](#) below. You can configure the client to publish ip-mac and ip-duid (for DHCPv6 leases) metadata at the Grid and member levels. You can also configure the client to publish metadata for specific leases by overriding the Grid or member publishing settings at the network (IPv4 and IPv6) or range (IPv4 only) level. The DHCP server sends updates to the IF-MAP server using the XML format and SOAP/HTTPS bindings specified in IF-MAP v1.1r5 and v2.0r26. The DHCP server supports the IF-MAP 2.0 protocol by default. You can also enable the support for IF-MAP 1.1, as described in [Configuring a Grid to Support IF-MAP](#) below.

When the DHCP server grants an IPv4 lease and sends the DHCPACK packet to the DHCP client, it updates the link in the IF-MAP server between the leased IP address and client MAC address with ip-mac metadata with the following attributes: start-time, end-time, and dhcp-server. The dhcp-server attribute contains the DHCP server hostname. The ip-mac metadata is attached to a link with:

- An ip-address identifier with the type attribute set to IPv4, a value attribute that contains the leased IP address, and the administrative-domain attribute set to the network view to which the IP address belongs.
- A mac-address identifier with a value attribute that contains the client MAC address. It does not have the administrative-domain attribute.

When the DHCP server grants an IPv6 lease and sends the Reply message to the DHCP client, it updates the link in the IF-MAP server between the leased IP address and client DHCP Unique Identifier (DUID) with ip-duid metadata that contains the following attributes: start-time, end-time, and dhcp-server. The dhcp-server attribute contains the DHCP server hostname. The ip-duid metadata is attached to a link with:

- An ip-address identifier with the type attribute set to IPv6, a value attribute that contains the leased IP address, and the administrative-domain attribute set to the network view to which the IP address belongs.
- A duid identifier with a value attribute that contains the client DUID. It does not have the administrative-domain attribute.

The Infoblox DHCP server also publishes data when an IPv4 or IPv6 lease changes. When a lease is released or when an active lease expires, the DHCP server sends a "publish delete" request to the IF-MAP server.

You can define how the IF-MAP server handles the existing ip-mac and ip-duid information before the DHCP client sends the next update. For example, you can specify the IF-MAP server to always delete existing ip-mac and ip-duid information before the next update. For information, see [Deleting Existing Data Before Publishing](#) below.

Following are the tasks to enable DHCP servers in a Grid to function as IF-MAP clients:

1. Enable IF-MAP in the Grid and specify the URL and port of the IF-MAP server, as described in the next section [Configuring a Grid to Support IF-MAP](#).
2. Optionally, enable the validation of the IF-MAP server certificate and import the CA certificate, as described in [Validating the IF-MAP Server Certificate](#) below.
3. Enable IF-MAP on each Grid member and specify an authentication method the member uses to connect to the IF-MAP server, as described in [Configuring Members as IF-MAP Clients](#) below.
4. Optionally, override publishing settings at the member, network, or range level, as described in [Overriding IF-MAP Publishing Settings](#) below.

You can also delete DHCP data published by a specific member, or define how the IF-MAP server deletes existing DHCP data before a client publishes an update. For information, see [Deleting Data from the IF-MAP Server](#) below.

### Configuring a Grid to Support IF-MAP

1. From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode**.
3. Click the **IF-MAP** tab, and then complete the following:

- **Enable IF-MAP:** Select this checkbox to enable the IF-MAP service for the Grid. Note that you must enable the IF-MAP service in order to enable or disable publishing at the Grid, member, network, or range level.
  - **IF-MAP Server URL:** Enter the URL of the IF-MAP server to which the Grid members publish DHCP data. The URL must begin with **https://**. For example, **https://<server\_ip\_addr>/ifmap**.
  - **IF-MAP Server Port:** The default HTTP port is 80 and the default HTTPS port is 443. Optionally, you can specify a different port on the IF-MAP server.
  - **Enable IF-MAP publishing:** Select this checkbox to enable IF-MAP publishing for the Grid. When you select this, IF-MAP publishing is enabled for all members, networks (IPv4 and IPv6), and DHCP ranges (IPv4 only). You can override the Grid property at a specific level to control the ip-mac and ip-uid metadata you want the client to publish for specific leases. For information, see [Overriding IF-MAP Publishing Settings](#) below.
  - **IF-MAP Protocol Version:** Select the IF-MAP protocol version you want the IF-MAP client to use to connect to the IF-MAP server. The default is IF-MAP 2.0.
4. Save the configuration and click **Restart** if it appears at the top of the screen.
  5. You can also configure how the IF-MAP server deletes existing metadata before the IF-MAP client publishes another update. For information, see [Deleting Data from the IF-MAP Server](#) below.

### Validating the IF-MAP Server Certificate

You can configure the IF-MAP client to validate the IF-MAP server certificate before the client establishes a connection or performs IF-MAP transactions. To validate an IF-MAP server certificate, you must first import the certificate of the CA that signs the IF-MAP server certificate.

To configure the IF-MAP client to validate the IF-MAP server certificate:

1. From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode**.
3. Click the **IF-MAP** tab and complete the following:
  - **Enable IF-MAP:** Select this checkbox to enable the IF-MAP service for the Grid.
  - **Enable IF-MAP server certificate validation:** Select this checkbox to enable the validation of the IF-MAP server certificate, and then click **Import** to import the CA certificate. In the *Upload* dialog box, click **Select** to navigate to the certificate, and then click **Upload**. You can also copy and paste the CA certificate here.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

### Configuring Members as IF-MAP Clients

To configure a member to be an IF-MAP client, you must first enable IF-MAP on the member and then configure a client authentication method. The IF-MAP client can authenticate itself to the IF-MAP server through user name and password credentials or digital certificate. Note that each member must have unique credentials or certificates. You cannot use the same credentials or certificates on multiple members. The appliance supports only one CA-signed certificate on each member. If you want to use a roll-over certificate, you must replace the existing certificate and restart services on the member.

To enable an appliance to function as an IF-MAP client:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.
2. In the *Member DHCP Properties* dialog box, click **Toggle Advanced Mode**.
3. Click the **IF-MAP** tab and complete the following:
  - **Enable IF-MAP:** Select this checkbox to enable the IF-MAP service on the member. Note that you must enable the IF-MAP service in order to enable or disable publishing at the network and range levels.
  - **Authentication:** Select one of the following authentication methods:
    - **Certificate:** Select this to use the IF-MAP client certificate for client authentication. You must already have a certificate configured for the member before you can select and save this configuration. For information about creating a client certificate, see [Creating IF-MAP Client Certificates](#) below.
    - **Basic:** Select this to use username and password credentials for IF-MAP client authentication. Complete the following:

- **Username:** Enter the username the member uses to connect to the IF-MAP server. This username must have been configured as a valid username on the IF-MAP server. Each member must have its own username.
  - **Password:** Enter the password the member uses to connect to the IF-MAP server.
  - **ConfirmPassword:** Enter the password again.  
Note when you upgrade to a new NIOS release, the basic authentication credentials are retained if IF-MAP was enabled and basic authentication was used before the upgrade.
  - **Enable IF-MAP publishing:** Click **Override** to override the Grid setting. Select this checkbox to enable IF-MAP publishing for all the networks that are served by this member. Ensure that you enable IF-MAP at either the Grid or member level in order to enable IF-MAP publishing for all networks.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Creating IF-MAP Client Certificates

Before you can select "Certificate" as the client authentication method, you must first create a certificate for the specified member.

You can do one of the following to generate an IF-MAP client certificate:

- Generate a self-signed certificate and save it. For information, see [Generating Self-Signed Certificates](#) below.
- Request a CA (Certificate Authority) signed certificate. When you receive the certificate from the CA, upload it to the member that you configure as an IF-MAP client.

## Generating Self-Signed Certificates

You can replace the default certificate with a self-signed certificate that you generate. When you generate a self-signed certificate, you can specify the correct hostname and change the public/private key size, enter valid dates and specify additional information specific to the member. If you have multiple members, you can generate a certificate for each appliance with the appropriate hostname.

To generate a self-signed certificate:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox, and then click **IF-MAP Client Certificate** -> **Generate Self-signed Certificate** from the Toolbar.
2. In the *Generate Self-Signed Certificate* dialog box, complete the following:
  - **Secure Hash Algorithm and Key Size:** You can select SHA-1 and a RSA key size of 1024 or 2048. SHA-256 (SHA-2) can be selected together with a RSA key size of 2048 or 4096. The default value is SHA-256 2048.
  - **Days Valid:** Specify the validity period of the certificate.
  - **Common Name:** Specify the domain name of the member. You can enter the FQDN (fully qualified domain name) of the appliance.
  - **Organization:** Enter the name of your company.
  - **Organizational Unit:** Enter the name of your department.
  - **Locality:** Enter a location, such as the city or town of your company.
  - **State or Province:** Enter the state or province.
  - **Country Code:** Enter the two-letter code that identifies the country, such as US.
  - **Admin E-mail Address:** Enter the email address of the appliance administrator.
  - **Comment:** Enter information about the certificate.
3. Click **OK**.
4. If the appliance already has an existing client certificate, the new certificate replaces the existing one. In the *Replace IFMAP Certificate Confirmation* dialog box, click **Yes**.

## Generating Certificate Signing Requests

You can generate a CSR (certificate signing request) that you use to obtain a signed certificate from your own trusted CA. Once you receive the signed certificate, you can import it to the member, as described in [Uploading Certificates](#) below.

To generate a CSR:



1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox, and then click **IF-MAP Client Certificate** -> **Create Signing Request** from the Toolbar.
2. In the *Create Certificate Signing Request* dialog box, enter the following:
  - **Secure Hash Algorithm and Key Size:** You can select SHA-1 and a RSA key size of 1024 or 2048. SHA-256 (SHA-2) can be selected together with a RSA key size of 2048 or 4096. The default value is SHA-256 2048.
  - **Common Name:** Specify the domain name of the member. You can enter the FQDN of the appliance.
  - **Organization:** Enter the name of your company.
  - **Organizational Unit:** Enter the name of your department.
  - **Locality:** Enter a location, such as the city or town of your company.
  - **State or Province:** Enter the state or province.
  - **Country Code:** Enter the two-letter code that identifies the country, such as US.
  - **Admin E-mail Address:** Enter the email address of the appliance administrator.
  - **Comment:** Enter information about the certificate.
3. Click **OK**.

### Uploading Certificates

When you receive the certificate from the CA, the appliance finds the matching CSR and takes the private key associated with the CSR and associates it with the newly imported certificate. The appliance then automatically deletes the CSR. To import a certificate:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox, and then click **IF-MAP Client Certificate** -> **Upload Certificate** from the Toolbar.
2. Navigate to where the certificate is located and click **Open**.
3. If the appliance already has an existing IF-MAP client certificate, the new certificate replaces the existing one. In the *Replace IF-MAP Certificate Confirmation* dialog box, click **Yes**.

### Downloading Certificates

You can download the current certificate or a self-signed certificate. To download a certificate:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox, and then click **IF-MAP Client Certificate** -> **Download Certificate** from the Toolbar.
2. Navigate to where you want to save the certificate, enter the file name, and then click **Save**.

### Overriding IF-MAP Publishing Settings

When you enable IF-MAP publishing at the Grid level, all members, networks (IPv4 and IPv6), and DHCP ranges (IPv4 only) in the Grid inherit the same setting. To control which ip-mac and ip-duid metadata is published for specific leases that belong to a specific network or address range, you can override the Grid settings at a specific member, network, or range level. Note that you must first enable the IF-MAP service at the Grid and member levels in order to enable or disable IF-MAP publishing at other levels. For example, if you want the DHCP server to publish IF-MAP data for specific leases in a specific network that is served by a specific member, you must first enable the IF-MAP service at the Grid and member levels, as described in *Configuring a Grid to Support IF-MAP* above. Then, you can enable IF-MAP publishing at the range level, as described in this section.

Though you can configure and save the settings of IF-MAP publishing any time at any level, the publishing does not actually happen unless the IF-MAP service is enabled at the Grid or member level. If a network or DHCP range is served by a specific member and you want to enable IF-MAP publishing for the network or range, you must first enable the IF-MAP service for the specified member.

To override IF-MAP publishing settings:

1. **Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.  
**Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> *network* checkbox, and then click the Edit icon.  
**DHCP Range:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> *network* -> *addr\_range* checkbox, and then click the Edit icon.
2. In the editor, click **Toggle Advanced Mode**, and then click the **IF-MAP** tab.

3. Click **Override** and complete the following:
  - **Enable IF-MAP Publishing:** Select this checkbox to instruct the DHCP server to publish metadata to the IF-MAP server when the IF-MAP service is enabled for the Grid or member. Clear this checkbox so the DHCP server does not publish metadata to the server.

### Deleting Data from the IF-MAP Server

The appliance allows you to delete IF-MAP data from the IF-MAP server. You can delete all IF-MAP data published by a specific member. You can also define how the IF-MAP server handles the deletion of existing metadata before the IF-MAP client publishes another update.

### Deleting All Data

You can delete all IF-MAP data published by a specified member. To delete data published by all members in a Grid, you must delete data for each member individually.

To delete IF-MAP data published by a member from the IF-MAP database:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab, and then click **Clear** -> **IF-MAP Data** from the Toolbar.
2. In the *Purge IF-MAP Data* dialog box, click **Select Member** to select a member. If there are multiple members, Grid Manager displays the *Member Selector* dialog box from which you can select one. Click the member name in the dialog box, and then click **Purge** to delete all the DHCP data published by the Grid member. You can also click **Clear** to clear the displayed member and select a new one.

### Deleting Existing Data Before Publishing

You can define how the IF-MAP server deletes existing metadata before an IF-MAP client publishes new data. You can configure the IF-MAP client to instruct the server to always delete existing data, never delete it, or delete the data before a specified time period.

To define how the IF-MAP server deletes DHCP data before the next publish:

1. From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode**.
3. Click the **IF-MAP** tab and complete the following:
  - **Enable IF-MAP:** Select this checkbox to enable the IF-MAP service.
  - **Delete existing metadata:** You can define how the IF-MAP server deletes the existing metadata before the IF-MAP client publishes new data. Select one of the following:
    - **Always delete:** Select this to always delete existing metadata before the IF-MAP client publishes updates. This is the default.
    - **Do not delete:** Select this to never delete the existing metadata before the IF-MAP client publishes updates.
    - **Earlier than:** Select this to delete metadata that was published before a given time before the IF-MAP client publishes updates. When you select this option, enter a time value, and then select a time unit from the drop-down list.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

### Starting DHCP Services on a Member

The DHCP service is disabled by default. After you complete the DHCP configuration, you can start DHCP service on a member. To enable the member to provide DHCPv6 service as well, you must start the DHCP service and then enable the DHCPv6 service on the member. In addition, you must specify the DHCP Unique Identifier (DUID) of the member. IPv6 clients use DUIDs to identify the source of the DHCP messages from servers.

To start DHCP service on a member:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox.
2. Expand the Toolbar and click **Start**.
3. In the *Start Member DHCP Service* dialog box, click **Yes**.
4. Grid Manager starts DHCP on the selected member.

You can stop DHCP service on a member by selecting the member checkbox and click **Stop** from the Toolbar. This will stop DHCP service enabled on the LAN port.

To stop DHCP service enabled on the LAN2 port:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox.
2. Click the Edit icon.
3. In the *Member DHCP Properties* editor, select the **General Basic** tab.
4. Clear the checkbox for LAN2 under DHCP interfaces.
5. Save the configuration.

To enable DHCPv6 service on the member:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox.
2. In the *Member DHCP Properties* editor, select the **General Basic** tab.
3. In the **IPv6 Properties** section, do the following:
  - **Server DUID:** Enter the DUID of the member.
  - **Enable DHCPv6 Service:** Select this checkbox.
4. Save the configuration.

## Viewing DHCP Member Status

You can view DHCP member status after you configure DHCP properties and start or stop DHCP services on a member.

To view member status:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** section.
2. Grid Manager displays the following information:
  - **Name:** The name of the Grid member.
  - **Status:** The status of the DHCP services on the member. This can be one of the following:
    - **Not Running:** DHCP services have not been started on the member.
    - **Running:** The DHCP services are running properly on the member.
    - **Warning:** The member is connecting or synchronizing with its Grid Master.
    - **Error:** The member is offline, is not licensed (that is, it does not have a DNSone license with the Grid upgrade that permits Grid membership), is upgrading or downgrading, or is shutting down. Note that you can mouse over on the informational icon next to the status to view detailed information.
  - **Comment:** The information you entered for the member.
  - **IPv4 DHCP Utilization:** The percentage of the total IPv4 DHCP utilization of the member. This is the percentage of the total number of DHCP hosts, fixed addresses, reservations, and leases assigned to the member versus the total number of IP addresses (excluding IP addresses in the exclusion range) and all DHCP objects assigned to the member. Note that only enabled objects are included in the calculation. It does not include abandoned addresses or leases. The appliance updates the utilization data every 15 minutes. The appliance displays the utilization data in one of the following colors:
    - **Red:** The DHCP resources are 100% utilized.
    - **Yellow:** The utilization percentage is over the effective high watermark threshold.
    - **Blue:** The utilization percentage is below the effective low watermark threshold.
    - **Black:** The utilization percentage is at any number other than 100%, or within the effective thresholds.
  - **Site:** The site to which the member belongs. This is one of the predefined extensible attributes.

You can select the following additional columns for display:

- **Address:** The IP address of the member.
- **Static Addresses:** The number of static IP addresses.
- **Dynamic Addresses:** The number of dynamically assigned IP addresses.
- **IF-MAP Connection:** The status of the IF-MAP service connection on the member. This can be one of the following.
  - **Stopped:** The DHCP or IF-MAP service on the member is stopped, or the IF-MAP service is not enabled.
  - **Running:** The IF-MAP client is connected to the IF-MAP server and the IF-MAP service is running properly.
  - **Failed:** The IF-MAP client cannot publish data to the IF-MAP server due to some errors.

- **Warning:** Some non-fatal errors occurred. The IF-MAP client attempts to reconnect to the server.



#### Note

You can mouse over on the informational icon next to the status to view detailed information, including the status description and the timestamp when the status initially changed.

- **IF-MAP Last Update:** The timestamp the status of the IF-MAP service was last updated. For example, if the IF-MAP connection status is **Running** and this field shows 2011-11-20 12:30:42 EST, it means that an IF-MAP operation, such as a publish, was last completed on November 20, 2011 at 12:30:42 Eastern Standard Time.

To view status information about the IF-MAP connection on an independent appliance, from the **Data Management** tab -> **DHCP** tab, click **System DHCP Properties** from the toolbar. The appliance displays the following:

- **IF-MAP Connection:** The status of the IF-MAP service on the independent appliance. A color icon associated with the connection status appears before the status.
- **IF-MAP Connection Information:** Detailed information about the status. On a Grid member, this information appears when you mouse over on the informational icon.
- **IF-MAP Last Update:** The timestamp when the status of the IF-MAP service last changed.



#### Note

For more information about these fields, see descriptions about Grid member status in this section.

You can view detailed information about a specific member by clicking the member link. Grid Manager displays the following information about the selected member:

- **Network:** The network assigned to the member.
- **Comment:** The information about the network.
- **IPv4 DHCP Utilization:** The percentage of the DHCP usage of the network. This is the percentage of the total number of fixed addresses, reservations, hosts, and active leases on the network over the total IP addresses in the range, excluding the number of addresses on the network. Note that only enabled objects are included in the calculation. It does not include abandoned addresses or leases.
- **Site:** The site to which the DHCP object belongs. This is one of the predefined extensible attributes. In the member panel, you can select the following additional fields for display:
- **Disabled:** Indicates whether the member is disabled or not.
- **IPAM Utilization:** When you define a network, this is the percentage based on the IP addresses in use divided by the total addresses in the network. For example, in a /24 network, if there are 25 static IP addresses defined and a DHCP range that includes 100 addresses, the total number of IP addresses in use is 125. Of the possible 256 addresses in the network, the IPAM utilization is about 50% for this network.

When you define a network container that contains subnets, this is the percentage of the total address space defined within the container regardless of whether any of the IP addresses in the subnets are in use. For example, when you define a /16 network and then 64 /24 networks underneath it, the /16 network container is considered 25% utilized even when none of the IP addresses in the /24 networks is in use.

You can use this information to verify if there is a sufficient number of available addresses in a network. The IPAM utilization is calculated approximately every 15 minutes.

- Extensible attributes that associate with the network.

You can also sort the data in ascending or descending order by column. For information, see [Customizing Tables](#).

## Viewing DHCP Configuration Files

You can view the IPv4 and IPv6 DHCP configuration of a selected member. The format of the configuration file depends on the browser you use.

To view the DHCP configuration of a selected member:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox.
2. Expand the Toolbar, select **View DHCP Configuration**, and then select either **IPv4** or **IPv6**. Grid Manager displays the IPv4 or IPv6 DHCP configuration of the selected member in a new browser. You can print and save the file using the corresponding functions in your browser.

## Managing IPv4 DHCP Data

To configure DHCP service for an IPv4 network and the resources in the network, perform the following tasks:

1. Create a network and assign it to Grid members or Microsoft DHCP servers, or an interface on an external discovered device. For information about Adding IPv4 Networks and Modifying IPv4 Networks, see [Configuring IPv4 Networks](#).
2. Configure DHCP properties for the network. You can override properties set at the Grid or member level and enter unique values for the network. For information, see [Configuring General IPv4 DHCP Properties](#) and [Configuring DHCP IPv4 and IPv6 Common Properties](#).
3. Optionally, assign zones to a network. For information, see [Associating Networks with Zones](#).
4. Add a DHCP range to the network and assign it to a member, a failover association, or a Microsoft DHCP server. For information about Adding IPv4 Address Ranges and Modifying IPv4 Address Ranges, see [Configuring IPv4 Address Ranges](#).
5. Optionally, add exclusions to the DHCP range for addresses that are not used for dynamic allocation. For information, see [Configuring IPv4 Fixed Addresses](#).
6. Optionally, configure DHCP properties for the address range. You can override properties set at an upper level and enter unique values for the address range. For information, see [Modifying IPv4 Address Ranges](#).
7. Optionally, define filters for precise address assignments and apply them to the DHCP range. For information, see [About IPv4 DHCP Filters](#).
8. Optionally, add fixed addresses and reservations to the network and configure DHCP properties for them. For information, see [Configuring IPv4 Fixed Addresses](#) and [Configuring IPv4 Reservations](#).

This section explains how to configure and manage IPv4 DHCP data. It contains the following topics:

- [Configuring the Next Available Network or IP Address](#)
- [Configuring IPv4 Networks](#)
- [Configuring IPv4 Shared Networks](#)
- [Configuring IPv4 Address Ranges](#)
- [Configuring IPv4 Fixed Addresses](#)
- [Configuring IPv4 Reservations](#)
- [Viewing IPv4 DHCP Objects](#)
- [About Roaming Hosts](#)

### Configuring the Next Available Network or IP Address

When you create certain objects through Grid Manager, the appliance can obtain the next available IPv4 or IPv6 network from a specific network container. It can also obtain the next available IP address from a specific network or address range. This feature automates the allocation of networks and IP addresses so you can manage your network space more efficiently. You can also use this feature to organize network devices. For example, you can create a reserved range called "Printer Range" to reserve static IP addresses for printers in your network. When you allocate IP addresses for printers, you can have the appliance search for the next available IP address within "Printer Range," and then allocate the next available address to a new printer.

When you create a new network, the appliance can look up the next available network address within a specific network container. The next available network address is the first unused network address in the network container to which you have administrative permissions. For information about creating IPv4 and IPv6 networks using the next available feature, see [Adding IPv4 Networks](#) and [Adding IPv6 Networks](#).

You can also obtain the next available IP address when you define a fixed address, reservation, or host record. The next available IP address is the first unused IP address in a specified network, DHCP range, or reserved range to which you have administrative permissions. For information about creating fixed addresses, reservations, and host records using

the next available feature, see [Configuring IPv4 Fixed Addresses](#), [Configuring IPv4 Reservations](#), and [Adding Host Records](#).

## Obtaining the Next Available

The appliance searches for the next available network or IP address based on the context you define when you create an object. For example, when you create a network within a specific network container, the appliance searches for the next available network within the specified container. Similarly, when you drill down to an address range and create an object from there, the appliance looks up the next available IP address within that address range. If you are not within a specific network or address range when you create an object, Grid Manager displays a selector from which you can select the network or address range for the next available network or IP address.

For information about how the appliance select the next available network and IP address, see [Guidelines for the Next Available Network and IP Address](#) below.

## Guidelines for the Next Available Network and IP Address

The appliance follows certain rules when searching for the next available network and IP address in the specified wizard, network container, and address range.

In a wizard where you can obtain the next available network or IP address, the following applies:

- In a wizard, if you add a network or IP address and then delete it, the appliance excludes it from the next available. When you try to obtain the next available network or IP address in the same wizard, the appliance does not return the deleted network or IP address until you exit the wizard.

In a network, the appliance searches for the next IP address that meets all of the following criteria:

- It does not match any DNS resource record, such as an A or PTR record, that is associated with an IP address.
- It is not assigned to a DHCP fixed address or host address record.
- It is not part of any DNS bulk host record.
- It does not match any unmanaged IP address.
- It is not the network (the first) or broadcast (the last) address in the specified network.
- It is not within any DHCP range in this network.
- It is not within any reserved range in this network.
- It is not within an exclusion range.
- It is not part of a scheduled task that involves a fixed address. For information about how to schedule a task, see [Scheduling Tasks](#).

In a DHCP range, the appliance searches for the next IP address that meets all of the following criteria:

- It is not assigned to a fixed address or host record.
- It does not match any unmanaged IP address.
- It is not part of an exclusion range within the DHCP range.
- It is not part of a scheduled task that involves a fixed address.
- It does not match any active DHCP lease.

In a reserved range, the appliance searches for the next IP address that meets all of the following criteria:

- It is not assigned to a fixed address or host record.
- It does not match any unmanaged IP address.
- It is not part of a scheduled task that involves a fixed address.



### Note

The appliance does not search for deleted leases in the Recycle Bin.

When multiple users simultaneously request for the next available network or IP address, the appliance returns the same unused network or IP address to all users. The user who first saves the task gets the next available network or IP address. In some cases, other users get an error message telling them that the network or IP address is not available when they save their tasks. They can then request another unused network or IP address or enter a new one.

## Configuring IPv4 Networks

When you create an IPv4 network, you can do so from scratch or from a network template that contains predefined properties. When you use a template to create a network, the properties of the template apply to the new network. For information about network templates, see [About IPv4 DHCP Templates](#). You can also create an IPv4 network from the Tasks Dashboard, as described in [Tasks Dashboard](#).

After you create IPv4 networks, you can combine them into shared networks or create ranges and fixed addresses.

### Adding IPv4 Networks

When you configure an IPv4 network, you must assign either Grid members or Microsoft servers to the network. A network cannot be served by a mix of Microsoft and Infoblox DHCP servers. Multiple servers can serve a network, but Grid members and Microsoft servers cannot serve the same network.

A Grid member can serve only one network view. Similarly, a Microsoft server can serve only one network view. Therefore when you assign Grid members to networks, you must assign the members to networks in the same network view. For information, see [Managing IPv4 DHCP Data](#).

If you have enabled support for RIR (Regional Internet Registry) updates and are adding an RIR IPv4 network container or network to NIOS, Grid Manager displays an RIR section in the *Add IPv4 Network* wizard. You must enter RIR related information in this section in order for NIOS to associate the newly added network with an RIR organization. For more information about RIR address allocation and updates, see [RIR Registration Updates](#).

To add an IPv4 network, complete the following:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab.
2. In the **Networks** section, expand the **Toolbar** and click the **Add** drop-down list and select **Network** -> **IPv4** network or, click the Add icon drop-down list and select **IPv4 Network**.
3. In the *Add Network* wizard, select one of the following and click **Next**:
  - **Add Network**: Click this to add a network from scratch.
  - **Add Network using Template**: To use a template, click this, and then click **Select Template** and select a network template. For information, see [About IPv4 DHCP Options](#). The appliance populates the template properties in the wizard when you click **Next**. You can then edit the pre-populated properties, except for **Netmask**.
4. Complete the following and then click **Next**:
  - **Regional Internet Registry**: This section appears only when support for RIR updates is enabled. For information about RIR, see [RIR Registration Updates](#). Complete the following to create an RIR IPv4 network container or network:
    - **Internet Registry**: Select the RIR from the drop-down list. The default is **RIPE**. When you select **None**, the network is not associated with an RIR organization.
    - **Organization ID**: Click **Select Organization** and select an organization from the *RIR Organization Selector* dialog box.
    - **Registration Status**: The default is **Not Registered**. When adding an RIR allocated network, you can change this to **Registered** and select the **Do not update registrations** checkbox below. Note that when you select **API** as the communication method, the registration status will be updated automatically after the registration update is completed. However, when you select **Email** as the communication method, the registration status will not be automatically updated. If you are creating a new network and the registration update is completed successfully, the status will be changed to **Registered**. If the update fails, the status will be changed to **Not Registered**. The updated status and timestamp are displayed in the **Status of last update** field in the *IPv4 Network Container* or *IPv4 Network* editor.
    - **Registration Action**: Select the registration action from the drop-down list. When you select **Create**, the appliance creates the IPv4 network and assigns it to the selected organization. When you select **None**, the appliance does not send registration updates to RIPE. When you are adding an existing RIR allocated network to NIOS, select **None**. When you are adding networks to an RIR allocated network (a parent network), select **Create**. Ensure that the parent network associated with an RIR organization already exists.



- **Do not update registrations:** Select this checkbox if you do not want the appliance to submit RIR updates to RIPE. By default, the appliance sends updates to the RIR database based on the configured communication method.
- **Network View:** This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the network.
- **Netmask:** Enter the netmask or use the netmask slider to select the appropriate number of subnet mask bits for the network. The appliance supports /1 to /32 netmasks. Note that when you use a template that contains a fixed netmask, you cannot adjust the netmask for this network.  
Microsoft servers can serve networks with /1 to /31 netmasks. Infoblox DHCP servers can serve networks with /8 to /32 netmasks.  
Since Infoblox DHCP servers do not support /1 to /7 networks, you can assign these networks to Microsoft DHCP servers only. You can create DHCP ranges and fixed addresses within these subnets.
- **Networks:** Do one of the following to add new networks:
  - Click the Add icon to enter a new network. Grid Manager adds a row to the table. Enter the network address in the **Network** field. Click the Add icon again to add another network.  
Or
  - Click the Next Available icon to have the appliance search for the next available network. Complete the following in the Next Available Networks section:
  - **Create new network(s) under:** Enter the network container in which you want to create the new network. When you enter a network that does not exist, the appliance adds it as a network container. When you enter a network that is part of a parent network, the parent network is converted into a network container if it does not have a member assignment or does not contain overlapping address ranges, fixed addresses, reservations, shared networks, and host records. When you enter a network that has a lower CIDR than an existing network, the appliance creates the network as a parent network and displays a message indicating that the newly created network overlaps an existing network. You can also click **Select Network** to select a specific network in the *Network Selector* dialog box. For information about how the appliance searches for the next available network, see [Obtaining the Next Available](#).
  - **Number of new networks:** Enter the number of networks you want to add to the selected network container. Note that if there is not enough network space in the selected network to create the number of networks specified here, Grid Manager displays an error message. The maximum number is 20 at a time. Note that when you have existing networks in the table and you select one, the number you enter here includes the selected network.
  - Click **Add Next** to add the networks. Grid Manager lists the networks in the table. You can click **Cancel** to reset the values.  
Note that you must click Add Next to add the network container you enter in the Next Available Networks section. If you enter a network in the Next Available Networks section and then use the Add icon to add another network, the appliance does not save the network you enter in the Next Available Networks section until you click Add Next.
- **Comment:** Enter useful information about the network, such as the name of the organization it serves.
- The **Sync to MGM** drop-down list is available only on the managed Grid when it remains joined with the Multi-Grid Master. Select one of the following from the **Sync to MGM** drop-down list:
  - **Yes:** Select this to enable synchronization of networks between the managed Grid and Multi-Grid Master.
  - **No:** Select this to disable synchronization of networks between the managed Grid and Multi-Grid Master.  
Note that if you have selected No at the parent network (disabled synchronization) and if you try to select Yes when adding a child network, the appliance returns an error. This means that you cannot override the settings at the child level if you have already restricted synchronization at the parent network.
  - **Use Inherited Setting:** Select this is to inherit synchronization settings from the parent object.
- **Automatically Create Reverse-Mapping Zone:** This function is enabled if the netmask of the network equals /8, /16, or /24. Select this to have the appliance automatically create reverse-mapping zones for the network. A reverse-mapping zone is an area of network space for which one or more name servers have the responsibility for responding to address-to-name queries. These zones are created in the DNS view assigned to receive dynamic DNS updates at the network view level.
- **Disable for DHCP:** Select this if you do not want the DHCP server to provide DHCP services for this network at this time. This feature is useful when you are in the process of setting up the DHCP server. Clear this after you have configured the server and are ready to have it serve DHCP for this network. Note

that disabling an IPv4 network may take a longer time to complete depending on the size of the data.

The **Cloud** section appears when the Cloud Network Automation license is installed on the Grid Master. For information, see [Deploying Cloud Network Automation](#).

To delegate authority for this network, complete the following:

#### Delegate authority from the Grid Master

- **Delegate To:** This field indicates whether the authority for the network you want to create has already been delegated to a Cloud Platform Appliance. Click **Select** to choose the Cloud Platform Appliance to which you want to delegate authority. The *Member Selector* displays only Cloud Platform Appliances in the Grid. Click the member, and Grid Manager displays the member name next to this field. This cloud member now assumes authority for this network, and the Grid Master does not have authority anymore. You can also click **Clear** to remove authority delegation from the selected Cloud Platform Appliance and return authority back to the Grid Master.
5. Click **Next** and add a Grid member or Microsoft server as a DHCP server for the network. A network can be served by either Grid members or Microsoft servers, but not both at the same time.
    - Click the Add icon and select one of the following options:
      - **Add Infoblox Member:** Select this option to add a Grid member as a DHCP server for the network. Select the Grid member from the *Member Selector* dialog box. Keep in mind, DHCP properties for the network are inherited from this member. The network can be served by multiple members, but a member can serve networks in one network view only.  
Or
      - **Add Microsoft Server:** Select this option to add a Microsoft server as a DHCP server for the network. Select the Microsoft server from the *Microsoft Server Selector* dialog box.
  6. Click **Next** to associate Active Directory Sites with the network. For more information, see [Associating Active Directory Sites with Networks](#).
  7. Click **Next** to override DHCP properties as described in [Configuring DHCP Properties](#). This only applies if you are adding a network that is served by an Infoblox Grid member.  
Or
  8. (*Applies only with Network Insight*) Click **Next** to initiate or disable discovery of the new network(s). *This step is not required for creating a new network.* Discovery settings differ based on whether you are defining one network or multiple networks.
    - **Configuring one network:** Discovery settings include the following: **Enable Discovery** and **Enable Immediate Discovery**, selecting a Probe member to perform the discovery; and **Polling Options**, which define how the network will be discovered by the Probe member, including the ability to enable or disable the use of **SNMP Credentials** and **CLI Credentials**, along with **Switch Port Data Collection** settings. By default, all Polling Options discovery settings are inherited from the parent network (or Grid, if no parent exists) unless you click **Override**. Polling Options govern the protocols used to query and collect information about the network devices being discovered. For more information, see the section [Configuring Discovery Properties](#) for a complete description of discovery polling options.  
Or
    - **Configuring more than one network:** If the networks are child networks, they automatically inherit the settings of the parent network, including discovery settings and the discovery member. Discovery is disabled for any parent networks. These settings will not appear in the wizard page. For discovery of multiple networks, you can only enable or disable **Immediate Discovery**.
  9. Assign VLAN objects to the network. For more information, see [VLAN Management](#).
  10. As part of creating a network in IPAM or DHCP, you can provision the network on an actual device (switch, router, or switch-router), that is discovered and managed through the Grid Manager.
    - Begin by checking the **Enable Network Provisioning** checkbox, and clicking the **Select Device** button. Choose your device from the Device Selector dialog. Click **Clear** to remove the setting. For more information, see the section [Using the Device Selector](#).
    - If you performed DHCP configuration in the previous step of the Add Network Wizard, the **Router IP** value will automatically be populated with the DHCP Router IP address value. Otherwise, you enter the standard router IP address.
    - For example, if Grid Manager discovers and manages a router 172.16.22.1, the IP value 172.16.22.1 should be entered in the **Router IP** field.

- If required for the newly provisioned network to ensure that attached devices receive DHCP auto-configuration, enable the **DHCP Forwarding** checkbox. For this setting, if a DHCP Failover was previously configured, the IP addresses defined for DHCP failover are automatically used for the DHCP forwarding configuration.
  - You will also need to choose an interface on the selected device on which to provision the network by selecting it from the **Interface** drop-down menu. Grid Manager ensures that only those interfaces that can support provisioning, and are available for provisioning (that do not have an **Operation Status** of Up), appear in the drop-down menu.
  - Otherwise, when creating networks and provisioning them on managed devices, you can create a VLAN on which to provision the network by clicking the **Create VLAN** option and entering the **VLAN Name** and **VLAN ID**. Ensure that the **VLAN ID** value you enter is appropriate for the application. Do not create a new VLAN and provision a network for a VLAN value that is already actively carrying traffic for another routing domain.  
If a selected device does not support VLANs, the **Create VLAN** option will not appear.
11. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Managing Extensible Attributes](#).  
If you are adding an RIR network, the RIR network attribute table appears. For information about these attributes and how to enter them, see [RIR Network Attributes](#). You can preview the information before the appliance submits updates to the RIPE database. To preview registration updates, click **Preview RIR Submissions**. For more information, see [Previewing Registration Updates](#).  
You cannot leave an optional RIR attribute value empty. If you do not have a value for an RIR attribute, you must delete it from the table. You can enter up to 256 characters for all RIR attributes.  
You need to assign a **Subscriber Member Site** to add subscriber service related extensible attributes in order to populate Subscriber Cache.
  12. As the final step in the Add IPv4 Network wizard, you define when Grid Manager provisions the new network by scheduling it. You also schedule when the associated port control task executes (if a port configuration has been specified).
    - To create the new network and its associated port configuration immediately, select **Now**. Grid Manager synchronizes the port control task to take place at the same time as the creation of the new network.
    - You can choose to have Grid Manager execute the port control task at a later time by selecting **Later**.
    - Choose a **Selected time** by entering or selecting a Start Date (click the calendar icon to choose a calendar date) and a Start Time, and choose a Time Zone.
  13. Choose one of the following from the **Save & ...** drop-down menu:
    - Click **Save & Close** to add the new network and close the wizard (this is the default).
    - Click **Save & Edit** to add the new network and launch the editor.
    - Click **Save & New** to add the new network and launch the wizard again to add another network.



#### Note

At any step during the wizard, you can click **Schedule for Later** to schedule the task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

After you create a network, you can perform the following:

- Use the split network feature to create subnets for the network. For information, see [Splitting IPv4 Networks into Subnets](#).
- Use the join networks feature to create a parent network that encompasses multiple subnets into a larger network. For information, see [Joining IPv4 Networks](#). You can also create a shared network for subnets that are on the same network segment.

Networks served by Microsoft servers do not support the split and join functions.

## Viewing Networks

You can view IPv4 networks from the **IPAM** tab -> Net Map and List panels. The Net Map panel provides a graphical view of your networks and the List panel displays the networks in table format. For more information about *Pv4 Network Map* and *IPAM Home* see, [Managing IPv4 Networks](#).

You can also view a list of IPv4 and IPv6 networks in the **DHCP** tab -> **Networks** tab -> **Networks** panel. This panel displays all IPv4 and IPv6 networks.

In any of these panels, you can use filters or the **Go to** function to navigate to a specific network. You can also create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#). You can add, delete, or edit a network. You can also monitor the DHCP utilization of a selected network.

When viewing networks, you can choose to view them in one of the following views:

- Click **Toggle Hierarchical View** to view networks hierarchically in a parent-child structure (networks in a network container). You can view detailed information about the networks by clicking the network link and drilling down to the lowest hierarchical level, and then opening a network. To go back to a previous hierarchical view, click the link of the corresponding level in the breadcrumb. The hierarchical view is the default view.
- Click **Toggle Flat View** to display a flat list of all networks and network containers. The parent and child networks are listed separately in the flat view.

Depending on where you view your networks, Grid Manager displays some of the following information by default. You can also select specific information for display.

- **Network:** The network address.  
The network is displayed in one of the following colors:
  - **Yellow:** The network is unmanaged.
  - **Blue:** The selected network.
  - **Gray:** The network is currently not available as a NIOS network object.
- **Comment:** The information you entered about the network.
- **IPAM Utilization:** This information is available for IPv4 networks only. It displays the percentage based on the IP addresses in use divided by the total addresses in the network. For example, in a /24 network, if there are 25 static IP addresses defined and a DHCP range that includes 100 addresses, the total number of IP addresses in use is 125. Of the possible 256 addresses in the network, the IPAM utilization is about 50% for this network. The appliance updates the IPAM utilization data immediately for a network container, but for a network, it is updated every 15 minutes.  
The IPAM utilization data is displayed in one of the following colors:
  - **Red:** The IPAM utilization percentage is above the configured Trigger value.
  - **Blue:** The IPAM utilization percentage is below the configured Trigger value.
- **Site:** The site to which the network belongs. This is one of the predefined extensible attributes.
- **Protocol:** Displays whether the network is an IPv4 or IPv6 network.
- **Disabled:** Indicates if the network is disabled. You can double-click a row and select the checkbox in this column to disable the network. Grid Manager displays a warning message when you select the checkbox. Click **Yes** to confirm or **No** to cancel. Note that disabling an IPv4 network may take a longer time to complete depending on the size of the data.
- **Leaf Network:** Indicates whether the network is a leaf network or not. A leaf network is a network that does not contain other networks.
- **IPv4 DHCP Utilization:** This information is available for IPv4 networks only. It displays the percentage of the total DHCP usage of the IPv4 network. This is the percentage of the total number of DHCP hosts, fixed addresses, reservations, and active leases in the network divided by the total number of IP addresses (excluding IP addresses in the exclusion range) and all DHCP objects in the network. Note that only enabled addresses are included in the calculation. It does not include abandoned addresses or leases. The appliance updates the utilization data approximately every 15 minutes. The utilization data is displayed in one of the following colors:
  - **Red:** The DHCP resources are 100% utilized.
  - **Yellow:** The DHCP utilization percentage is over the effective high-water mark threshold.
  - **Blue:** The DHCP utilization percentage is below the effective low-water mark threshold.
  - **Black:** The DHCP utilization percentage is at any number other than 100%, or it is not above and below any threshold.

You see the following when RIR is enabled (for more information, see [RIR Registration Updates](#)):

- **RIR Organization:** This appears only if support for RIR updates is enabled. This displays the name of the RIR organization to which the network is assigned.
- **RIR Organization ID:** This appears only if support for RIR updates is enabled. This displays the ID of the RIR organization to which the network is assigned.

- **RIR Registration Status:** This appears only if support for RIR update is enabled. This field displays the RIR registration status. This can be **Registered** or **Not Registered**. **Registered** indicates that the network has a corresponding entry in the RIPE database.
- **Last Registration Updated:** Displays the timestamp when the last registration was updated. The displayed timestamp reflects the timestamp used on the Grid Master.
- **Status of Last Registration Update:** Displays the registration status and communication method of the last registration update. The status can be Pending, Sent, Succeeded, or Failed. Each time you send a registration update to create, modify, or delete a network container or network, the updated status will be displayed here. If you have selected not to send registration updates, the previous status is retained.

You see the following only with Network Insight (For information, see [Infoblox Network Insight](#)):

- **Discovery Enabled:** Indicates whether discovery is allowed on the network container or the network.
- **Managed:** Indicates whether the network is set to Managed status under NIOS.
- **First Discovered:** The date and timestamp of the first occasion that NIOS discovered the network.
- **Last Discovered:** The date and timestamp of the last occasion that NIOS performed discovery on the network.

You see the following when the Cloud Network Automation license is installed on the Grid Master (For information, see [Deploying Cloud Network Automation](#)):

- **Cloud Usage:** This field indicates whether this object is associated with any specific cloud extensible attributes or within the scope of delegation. It can be one of the following:
  - **Cloud from adapter:** Indicates that this object has been created by a cloud adapter and it may or may not be within a scope of delegation at the moment.
  - **Cloud from delegation:** Indicates that this object is within the scope of delegation or the object itself defines a scope of authority delegation, and it is not created by a cloud adapter.
  - **Used by cloud:** Indicates that this network or network container is associated with the extensible attribute **Is External** or **Is Shared** and the value is set to True, which implies the network is a private or shared network managed by the CMP, and it is not **Cloud from adapter** or **Cloud from delegation**.
  - **Non-cloud:** The object is a regular NIOS object and is not within the scope of any authority delegation nor is it associated with any of these extensible attributes: **Cloud API Owned**, **Is External**, or **Is Shared**. NIOS admin users can modify this object based on their permissions.
- **Owned By:** A cloud object can be owned by the Grid Master or the cloud adapter. When the object is created by the Cloud Platform member, this shows **Grid**. If the object is created by the cloud adapter, this shows **Adapter**.
- **Delegated To:** This tells you whether a cloud object has been delegated to a Cloud Platform Appliance or not. If the cloud object has a parent object and the parent has been delegated, this field shows the parent delegation and you cannot modify the field.
- Extensible attributes (**Building**, **Country**, **Region**, **State**, and **VLAN**): You can select the extensible attributes such as Building, Country, Region, State, and VLAN for display. When you enable other features such as RIR, Network Insight, and Cloud Network Automation, you can select additional attributes for display.

You can sort the list of networks in ascending or descending order by certain columns. For information about customizing tables in Grid Manager, see [Customizing Tables](#).

You can also modify some of the data in the table. Double-click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#). You can edit values of inheritable extensible attributes by double-clicking on the respective row. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you perform an inline editing. The *Descendant Actions* dialog box is displayed when you click **Save**. For information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#). If you delete the value of an inheritable extensible attribute at the parent level, you can choose to preserve the descendant value or remove it. For information, see [Deleting Inheritable Extensible Attributes Associated with Parent Objects](#).

## Viewing Network Details

You can view detailed information about a specific network by clicking the network link. Grid Manager displays the objects in the network, including DHCP ranges, hosts, fixed addresses and roaming hosts. It displays the following information about the network:

- **IP Address:** The IP address of a DHCP object, such as a DHCP range, fixed address, reservation, host configured for DHCP, or roaming host with an allocated IP address. For a DHCP range, this field displays the start



and end addresses of the range. For a host that has multiple IP addresses, each IP address is displayed separately. Note that the appliance highlights all disabled DHCP objects in gray.

The IP address is displayed in one of the following colors:

- **Yellow:** The IP address is unmanaged.
- **Blue:** The IP address is excluded.
- **Pink:** Indicates IP address conflicts.
- **Gray:** The IP address is currently not available as a NIOS object.
- **Type:** The DHCP object type, such as **DHCP Range** or **Fixed Address**.
- **Name:** The object name. For example, if the IP address belongs to a host record, this field displays the host name.
- **Comment:** The information you entered for the object.
- **IPv4 DHCP Utilization:** The percentage of the total DHCP usage of a DHCP range. This is the percentage of the total number of fixed addresses, reservations, hosts, and active leases in the DHCP range divided by the total IP addresses in the range, excluding the number of addresses in the exclusion ranges. Note that only enabled objects are included in the calculation. It does not include abandoned addresses or leases.
- **Site:** The site to which the DHCP object belongs. This is one of the predefined extensible attributes. You can select the following additional columns for display:
- **Static Addresses:** Indicates whether the IP address is a static address.
- **Dynamic Addresses:** Indicates whether the IP address is a dynamically assigned address.
- **Disabled:** Indicates whether the object is disabled. You can double-click a row and select the checkbox in this column to disable the network. Grid Manager displays a warning message when you select the checkbox. Click **Yes** to confirm or **No** to cancel. Note that disabling an IPv4 network may take a longer time to complete depending on the size of the data.
- **Priority:** Displays the priority of a DHCP range when NAC filters are applied.
- Available extensible attributes.

You can also perform the following in this panel:

- Modify some of the data in the table. Double-click a row, and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes. Note that some fields are read-only. You can edit values of inheritable extensible attributes by double-clicking on the respective row. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you perform inline editing. The *Descendant Actions* dialog box is displayed when you click **Save**. For information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#). If you delete the value of an inheritable extensible attribute at the parent level, you can choose to preserve the descendant value or remove it. For information, see [Deleting Inheritable Extensible Attributes Associated with Parent Objects](#).
- Sort the displayed data in ascending or descending order by column.
- Click **Go to IPAM View** to view information about the object in the **IPAM** tab.
- Add new objects, such as DHCP ranges, to the network.
- Delete or schedule the deletion of a selected object or multiple objects.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Print or export the data.

## Modifying IPv4 Networks

You can modify existing network settings and override the Grid or member DHCP properties, with the exception of the network address and netmask.

To modify an IPv4 network, perform the following:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* checkbox, and then click the Edit icon.  
Or  
From the **Data Management** tab, select the **IPAM** tab -> *network* checkbox, and then click the Edit icon.
2. The *IPv4 Network* editor contains the following basic tabs from which you can modify data:
  - **Genera Basic:** You can modify the following fields:
    - **Comment:** The information you entered for the network.

- **Disabled:** This field is displayed only if the selected network is a network without a child network under it. You can disable and enable existing networks instead of removing them from the database, if the selected network does not have a child subnet. This feature is especially helpful when you have to move or repair the server for a particular network. Note that disabling an IPv4 network may take a longer time to complete depending on the size of the data.

#### Restricting synchronization of a network

- **Disable sync to MGM:** Select this checkbox to disable synchronization of a network from the managed Grid to the Multi-Grid Master. This checkbox is available only on the managed Grid when it remains joined with the Multi-Grid Master.

When the Cloud Network Automation license is installed on the Grid Master, Grid Manager displays the following information in the **Cloud** section: **Cloud Usage**, **Owned By**, and **Delegated To**. You cannot modify these fields. For more information, see [Viewing Networks](#) below.

- **Member Assignment:** Add or delete a Grid member that provides DHCP services for this network. For information, see [Adding IPv4 Networks](#) above.
  - **IPv4 DHCP Options:** Keep the inherited DHCP properties or override them and enter unique settings for the network. For information, see [Defining IPv4 DHCP Options](#).
  - **Discovery:** Checking the **Enable Discovery** checkbox informs NIOS to begin discovering the network after you click **Save and Close**. You manage discovery polling settings local to each network from this page. For a complete overview of features on this page, see [Discovering Devices and Networks](#) and its subsections.
  - **Discovery Exclusions:** IP Addresses and IP ranges can be locally excluded from discovery by clicking the Add icon and selecting **Add IP Address** or **Add IP Range**. These IP addresses or IP ranges are selected from within the chosen network. For related information, see [Excluding IP Addresses from Discovery](#) and its subsections.
  - **Discovery Blackout:** Define extended time periods and regularly scheduled times when discovery and/or Port Configuration tasks will not take place on a network. Editing a network under DHCP, blackout settings apply only to the specified network. You also specify the scheduled time when the blackout period begins, and the duration of the blackout period. By default, the network inherits its discovery blackout settings from the Grid level. For related information, see [Defining Blackout Periods](#) and its subsections. Note that discovery blackout settings also can be defined for DHCP ranges.
  - **RIR Registration:** Modify RIR network information. This tab appears only when support for RIR updates is enabled. For information, see [Modifying RIR Network Data](#).
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of the extensible attributes. For information, see [About Extensible Attributes](#). You can edit values of inheritable extensible attributes by -clicking on the respective column. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you perform inline editing. The *Descendant Actions* dialog box is displayed when you click **Save**. For information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#). If you delete the value of an inheritable extensible attribute at the parent level, you can choose to preserve the descendant value or remove it. For information, see [Deleting Inheritable Extensible Attributes Associated with Parent Objects](#).
  - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#).
3. Optionally, click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
- **General Advanced:** You can associate zones with a network. For information, see [Associating Networks with Zones](#).
  - **IPv4 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the network. Note that you must click **Override** and select **Enable DDNS updates** for the DDNS settings you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).
  - **IPv4 BOOTP/PXE:** Keep the inherited BOOTP properties or override them and enter unique settings for the network. For information, see [Configuring IPv4 BOOTP and PXE Properties](#).
  - **Filters:** You can keep the inherited IPv4 logic filters or override them and add a new IPv4 logic filter. For information, see [Applying Filters to DHCP Objects](#).
  - **IPv4 DHCP Thresholds:** Keep the inherited thresholds settings or override them and enter unique settings for the network. For information, see [Configuring Thresholds for DHCP Ranges](#).



4. Save the configuration or click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Deleting IPv4 Networks

When you delete a network, all of its data, including all DHCP records, subnets, and records in its subnets, is deleted from the database. Due to the potentially large loss of data that can occur when you delete a network, the appliance stores the deleted network in the Recycle Bin. You can restore a deleted network from the Recycle Bin, if enabled. You can also disable a network instead of deleting it.

To delete a network, perform the following:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* checkbox, and then select **Delete** or **Schedule Delete** from the Delete drop-down menu.
2. To delete the network now, in the *Delete Confirmation* dialog box, click **Yes**. Grid Manager displays a warning message. Click **Yes** to continue or **No** to cancel the process. To schedule the deletion, see [Scheduling Deletions](#). You can also choose to export all network data to a CSV file before deleting. To do this, click the **Export & Delete** button. For more information about this option, see [Exporting and Deleting Networks](#).

The appliance puts the deleted network in the Recycle Bin, if enabled. Click **Restore** in the Recycle Bin to recover the deleted data. Click **Yes** in the *Restore Item* dialog box to restore or **No** to cancel the process.



### Note

- Any port reservation associated with the deleted network will also be deleted without user intervention.
- Deleting and restoring an IPv4 network may take a longer time to complete depending on the size of the data.
- You cannot delete a network that has a VLAN object assigned to it. For more information, see [Assigning VLANs to a Network](#).

## Exporting and Deleting Networks

Before deleting a network or scheduling a network for deletion, you can choose to export all network data to a CSV file. To do this, click the **DHCP** tab > **Networks** tab > **Networks** section > *network* checkbox, select **Delete** or **Schedule Delete** > **Export & Delete** button.

The deleted network CSV file is downloaded and saved to your system as a .zip file called DeletedNetworks.zip. All parent networks, sub-parent networks (if they exist), and DHCP objects are exported to the CSV file.

If you schedule a deletion, the **Export & Delete** option is available only for the **Delete Now** option. You cannot export network data if you schedule the deletion for a later date or time.

The **Export & Delete** option supports up to 1.5 million network objects in a single delete operation.

## Configuring IPv4 Shared Networks

You can combine individual contiguous networks into a shared network to allow the DHCP server to assign IP addresses from any subnet (that resides on the same network interface) in the shared network.

Before creating a shared network, you must first create the subnets. For example, you must first create the networks `10.32.1.0` and `10.30.0.0` before designating them to a shared network. For more information, see [About Shared Networks](#).

## Adding IPv4 Shared Networks

To add a shared network, complete the following:

1. Select the **Data Management** tab.
2. If you have more than one network view in the system, select the network view in which you want to add the network.
3. Select the **DHCP** tab -> **Networks** tab.
4. In the **Shared Networks** section, select **IPv4 Shared Network** from the Add icon drop-down menu, or expand the **Toolbar**, click the **Add** drop-down list and select **Shared Network -> IPv4**.
5. In the *Add IPv4 Shared Network* wizard, complete the following and click **Next**:
  - **Name**: Enter the name of the shared network.
  - **Comment**: Enter information about the shared network.
  - **Disabled**: Select this if you want to enable the shared network at a later time. You can disable and enable existing networks instead of removing them from the database. This feature is especially helpful when you have to move or repair the server for a particular network.
6. Perform the following to add networks:
  - a. Click the Add icon.
  - b. In the *Select Network* dialog box, select the networks that you want to include in the shared network. Ensure that the networks are served by the same Grid members to avoid DHCP inconsistencies.
7. Click **Next** to configure or override DHCP options as described in [Defining IPv4 DHCP Options](#).
8. Click **Next** to enter values for required extensible attributes or add optional extensible attributes for the shared network. For information, see [Using Extensible Attributes](#).
9. Save the configuration or click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Viewing Shared Networks

To view IPv4 and IPv6 shared networks, complete the following:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Shared Networks**.
2. Grid Manager displays the following information:
  - **Name**: The name of the shared network.
  - **Protocol**: Displays whether the network is an IPv4 or IPv6 network.
  - **Comment**: The information you entered about the shared network.
  - **IPv4 DHCP Utilization**: The percentage of the DHCP utilization of the networks that belong to the shared network. This is the percentage of the total number of available IP addresses from all the networks that belong to the shared network versus the total number of all IP addresses in all of the networks in the shared network.
  - **Site**: The site to which the shared network belongs. This is one of the predefined extensible attributes.

You can select **Disabled** or available extensible attributes for display. You also can view detailed information about a network in a shared network by clicking the network link.

In this panel, you can use filters or the **Go to** function to navigate to a specific network. You can also create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).

You can sort the list of networks in ascending or descending order by columns. For information about customizing tables in Grid Manager, see [Customizing Tables](#).

You can also modify some of the data in the table. Double-click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).

## Modifying IPv4 Shared Networks

You can modify existing network settings and override the Grid or member DHCP properties. To modify, complete the following:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Shared Networks** section -> *shared\_network* checkbox, and then click the Edit icon.
2. The *Shared Network* editor contains the following tabs from which you can modify data:
  - **General**: Modify the fields **Name**, **Comments**, and **Disabled** as described in the section Adding IPv4 Shared Networks above.
  - **Networks**: Displays the networks that are currently assigned to the shared network. You can add or delete a network. To add a network, click the Add icon. In the *Select Network* dialog box, select the network you want to add. To delete an existing network, select the *network* checkbox, and then click the Delete icon.
  - **Extensible Attributes**: Add and delete extensible attributes that are associated with a specific network. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions**: This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions Managing Permissions](#).
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data:
  - **IPv4 DHCP Options**: Keep the inherited DHCP properties or override them and enter unique settings for the shared network. For information, see [Defining IPv4 DHCP Options](#).
  - **IPv4 DDNS**: Keep the inherited DDNS settings or override them and enter unique settings for the shared network. Note that you must click **Override** and select **Enable DDN Supdates** for the DDNS settings you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).
  - **IPv4 BOOTP/PXE**: Keep the inherited BOOTP properties or override them and enter unique settings for the shared network. For information, see [Configuring IPv4 BOOTP and PXE Properties](#).
  - **Filters**: You can keep the inherited IPv4 logic filters or override them and add a new IPv4 logic filter. For information, see [Applying Filters to DHCP Objects](#).  
Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
4. Save the configuration and click **Restart** if it appears at the top of the screen.  
Or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Deleting IPv4 Shared Networks

Though you can delete the networks in a shared network, a shared network must have at least one network in it. To delete a shared network, complete the following:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Shared Networks** section -> *shared\_network* checkbox, and then select **Delete** or **Schedule Delete** from the drop-down menu.
2. To delete the shared network now, in the *Delete Confirmation* dialog box, click **Yes**. To schedule the deletion, see [Scheduling Deletions](#).

The appliance puts the deleted shared network in the Recycle Bin, if enabled.

## Configuring IPv4 Address Ranges

In a network, you define address ranges from which the DHCP server or failover association assigns IP addresses to client requests. When a DHCP client requests an IP address, the appliance allocates an address within a defined DHCP range. The DHCP client can use the assigned IP address until the lease expires.

When you do not assign a DHCP server or failover association to an address range, the range becomes a reserved range. A reserved range contains IP addresses that are reserved for static hosts, not for dynamic assignments. You can allocate the next available IP from a reserved range.

You can also apply filters to DHCP ranges to control how the DHCP server allocates IP addresses. For information about DHCP filters, see [Configuring DHCP Filters](#).

## Adding IPv4 Address Ranges

To add an IPv4 address range:

1. Navigate to the IPv4 network to which you want to add an address range, and then select **Range** from the Add drop down menu.  
or  
From any panel in the DHCP tab, expand the Toolbar and click **Add -> Range -> IPv4**.
  2. In the *Add IPv4 Range* wizard, select one of the following and click **Next**:
    - **Add Range**: Select this to add an address range from scratch.  
or
    - **Add Range Using Template**: Click **Select Template** and select the template that you want to use. Note that when you use a template to create an address range, the configurations of the template apply to the new range. The appliance automatically populates the range properties in the wizard. You can then edit the pre-populated properties.
  3. Complete the following:
    - **Network**: Click **Select Network**. Grid Manager displays the network address here if you have only one network configured. When there are multiple networks, Grid Manager displays the *Select Network* dialog box from which you can select one.
    - **Start**: Enter the first available IP address in the range.
    - **End**: Enter the last available IP address in the range.
    - **Name**: Optionally, enter a name for the range.
    - **Comment**: Enter additional information about the address range.
    - **Disabled for DHCP**: Select this if you want to save the configuration for the address range but do not want to activate the address range yet. You can clear this checkbox when you are ready to allocate addresses from this range.

The **Cloud** section appears when the Cloud Network Automation license is installed on the Grid Master. For information, see [Deploying Cloud Network Automation](#). To delegate authority for this range, complete the following:

**Delegate authority from the Grid Master**

    - **Delegate To**: This field indicates whether the authority for the range you want to create has already been delegated to a Cloud Platform Appliance. Click **Select** to choose the Cloud Platform Appliance to which you want to delegate authority. The *Member Selector* displays only Cloud Platform Appliances in the Grid. Click the member, and Grid Manager displays the member name next to this field. This cloud member now assumes authority for this range, and the Grid Master does not have authority any more. You can also click **Clear** to remove authority delegation from the selected Cloud Platform Appliance and return authority back to the Grid Master.
  4. Click **Next** and select one of the following:
    - **None (Reserved Range)**: Select this if you want to reserve this address range for static hosts. Addresses in this range cannot be allocated as dynamic addresses. You can allocate the next available IP from this range to a static host. This is selected by default.
    - **Grid Member**: Select this if you want a Grid member to serve DHCP for this address range. Select a Grid member from the drop-down list. The drop-down list displays only the Grid members that are associated with the network to which the DHCP range belongs.
    - **Failover Association**: Select this if you want a failover association to serve DHCP for this address range. Click **Select Association**. In the *DHCP Failover Association Selector* dialog box, choose a failover association, and then click the Select icon. The appliance lists failover associations that serve DHCP in the network view of the DHCP range. For information, see [Configuring DHCP Failover](#).
  5. Click **Next** to configure or override DHCP options as described in [Defining IPv4 DHCP Options](#).
- Steps 6-7 apply only in deployments using Network Insight discovery features. Otherwise, skip to Step 8.
6. Click **Next** to initiate or disable discovery of the new DHCP range.
  7. Discovery settings include the following: **Enable Discovery** and **Immediate Discovery**, selecting a Probe member to perform the discovery; and **Polling Options**, which define how the network will be discovered by the Probe member. By default, all Polling Options discovery settings are inherited from the parent network unless you click **Override**. Polling Options govern the protocols used to query and collect information about the network devices being discovered. For more information, see the section [Configuring Discovery Properties](#) for a complete description of discovery Polling Options.
  8. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).
  9. Save the configuration and click **Restart** if it appears at the top of the screen  
or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

For information on viewing address ranges in a network, see [Viewing IPv4 DHCP Objects](#).

## Modifying IPv4 Address Ranges

You can modify settings for the DHCP range. You can also define an exclusion range to prevent the appliance from assigning the addresses in the exclusion range to clients. IP addresses in an exclusion range are excluded from the pool of IP addresses. For more information, see [About Exclusion Ranges](#).

To modify an IPv4 address range:

1. From the **DataManagement** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *addr\_range* checkbox, and then click the Edit icon.

2. The *DHCPRange* editor contains the following basic tabs from which you can modify data:

- **General:** Modify the fields, except the network address, as described in [Adding IPv4 Address Ranges](#) above.

When the Cloud Network Automation license is installed on the Grid Master, Grid Manager displays the following information in the **Cloud** section: **Cloud Usage**, **Owned By**, and **Delegated To**. You cannot modify these fields. For more information, see [Adding IPv4 Address Ranges](#) above.

- **Member Assignment:** Modify the Grid member or failover association that provides DHCP services for the DHCP range as described in [Adding IPv4 Address Ranges](#) above.

When you change the member assignment of DHCP ranges from **Failover Association** to **Grid Member** and then back to **Failover Association**, the leases in the primary and secondary servers fall out of sync. To resynchronize the peers, the failover association of the secondary server is put in the Recover-Wait state for the entire duration of Maximum Client Lead Time while a forced recovery takes place. During this period, only the IP addresses allocated to the primary server are available for lease.

- **IPv4 DHCP Options:** Keep the inherited DHCP options or override them and enter unique settings for the DHCP range. For information, see [Defining IPv4 DHCP Options](#).
- **Extensible Attributes:** You can add and delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#). You can edit values of inheritable extensible attributes by double clicking on the respective column. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you perform an inline editing. The *Descendant Actions* dialog box is displayed when you click **Save**. For information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#). If you delete the value of an inheritable extensible attribute at the parent level, you can choose to preserve the descendant value or remove it. For information, see [Deleting Inheritable Extensible Attributes Associated with Parent Objects](#).
- **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions Managing Permissions](#).

3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.

- **IPv4 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the DHCP range. Note that you must click **Override** and select **Enable DDN Supdates** for the DDNS settings you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).
- **IPv4 BOOTP/PXE:** Keep the inherited BOOTP properties or override them and enter unique settings for the DHCP range. For information, see [Configuring IPv4 BOOTP and PXE Properties](#).
- **Exclusion Ranges:** Configure a range of IP addresses that the appliance does not use to assign to clients. You can use these exclusion addresses as static IP addresses. Enter the start and end addresses of the exclusion range, and optionally, enter information about this exclusion range.
- **IPv4 DHCP Thresholds:** Keep the inherited thresholds settings or override them and enter unique settings for the DHCP range. For information, see [Configuring Thresholds for DHCP Ranges](#).
- **Filters:** You can keep the inherited IPv4 logic filters or override them and add a new IPv4 logic filter. For information, see [Applying Filters to DHCP Objects](#).  
Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Controlling Lease Assignments

You can set parameters to control how the DHCP server responds to lease requests within a specific DHCP range. When you set a DHCP range to deny all leases requests, the appliance does not assign IP addresses within this range to DHCP clients. This is useful when you want DHCP clients with IP addresses within this range to obtain new IP addresses when they renew their leases. When a client with an IP address within this range broadcasts a DHCPREQUEST message for its old IP address, the authoritative DHCP server responds with a DHCPNAK. This causes the client to move to the INIT state and to send a DHCPDISCOVER message for a new IP address.

You can also configure the DHCP server to assign or deny IP addresses within a DHCP range to known and unknown DHCP clients. Known clients include roaming hosts and clients with fixed addresses or DHCP host entries. Unknown clients include clients that are not roaming hosts and clients that do not have fixed addresses or DHCP host entries. To control how the appliance assigns leases to client requests:

1. **DHCP Range:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> **network** -> *addr\_range* checkbox, and then click the Edit icon.
2. In the *IPv4 Range* editor, click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, click the **IPv4 DHCP Options** tab -> **Advanced** tab and complete the following:

The **IPv4 DHCP Options** tab is enabled when you select a **Grid Member** or **IPv4 DHCP Failover Association** in the **Member Assignment** tab.

- **Allow/Deny Clients**

- **Known Clients:** Select this checkbox, and then select **Allow** or **Deny** from the drop-down list to assign or deny IP addresses within this range to known DHCP clients. Known DHCP clients include roaming hosts and clients with fixed addresses or DHCP host entries. Note that the appliance cannot deny an IP address to a fixed address within this range. You must disable the fixed address if you do not want it to obtain an IP address here.
- **Unknown Clients:** Select this checkbox, and then select **Allow** or **Deny** from the drop-down list to assign or deny IP addresses within this range to unknown DHCP clients. Unknown DHCP clients include clients that are not roaming hosts and clients that do not have fixed addresses or DHCP host entries.
- **Deny Leases:** Select **Deny all lease requests for this range** to deny all lease requests from DHCP clients.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Deleting IPv4 Address Ranges

To delete a DHCP range:

- From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *addr\_range* checkbox, and then click the Delete icon.

## Configuring IPv4 Fixed Addresses

A fixed address represents a persistent link between an IP address and one of the following:

- MAC address
- Client identifier
- Circuit ID or remote ID in the DHCP relay agent option (option 82)

You can create fixed addresses as described in [Adding IPv4 Fixed Addresses](#) below or from the Tasks Dashboard. For information about the Tasks Dashboard, see [The Tasks Dashboard](#). You can also create a fixed address when you create a host record or when you convert an active, dynamically leased address to a fixed address. For more information, see [Adding Host Records](#) and [Converting DHCP Leases](#).



When you create a fixed address, you must define a host identifier that the DHCP server uses to match the DHCP client. Every time the DHCP client with the matching identifier requests an IP address, the DHCP server assigns it the same address.

When a DHCP client sends a DHCPDISCOVER, it can include the MAC address or a unique client identifier as option 61 in the DHCP section of the packet. Using a client identifier is especially useful for virtualized server processes that might be moved to different hardware platforms. For information about option 61, refer to *RFC2132, DHCP Options and BOOTP Vendor Extensions*. You can select either the MAC address or client identifier as the host identifier in a fixed address. The DHCP server matches the option 61 value in the client request using either the MAC address or client identifier, depending on your configuration. When a DHCP client renews an IP address using a matching MAC address or client identifier, the DHCP server tracks the allocation of IP addresses and reserves the same IP address for the client.

When you enter a MAC address, you can use one of the following formats:

- aa:bb:cc:dd:ee:ff — Six groups of two hexadecimal digits separated by colons (:)
- aa-bb-cc-dd-ee-ff — Six groups of two hexadecimal digits separated by hyphens (-)
- aabb.ccdd.eeff — Three groups of four hexadecimal digits separated by periods (.)
- aabbcc-ddeeff — Two groups of six hexadecimal digits separated by a hyphen (-)
- aabbccddeeff — One group of 12 hexadecimal digits without any separator

After you save the entry, the appliance displays the MAC address in the AA:BB:CC:DD:EE:FF format.

When a DHCP client requests an IP address through a DHCP relay agent, the agent adds either the circuit ID or remote ID or both, to the DHCP relay agent information option (option 82). For information, see [About the DHCP Relay Agent Option \(Option 82\)](#). When you select the DHCP relay agent option (circuit ID or remote ID) as the host identifier in a fixed address, the DHCP server matches the DHCP client request using either the circuit ID or the remote ID, depending on your configuration. When a DHCP client renews an IP address using a matching relay agent ID, the DHCP server tracks the allocation of IP addresses and reserves the same IP address for the client. Note that leases are not renewed at the standard renewal time (T1) when option 82 information is not available as a unicast renewal. Instead, leases are renewed at the rebinding time (T2) when renewals are sent as broadcasts to the relay agents and contain option 82 information. For information about how to configure the lease time, see [Configuring General IPv4 DHCP Properties](#).

## Adding IPv4 Fixed Addresses

For Cloud Network Automation, you can create IPv4 fixed addresses within the delegation authority of a Cloud Platform Appliance. The newly created fixed address is forwarded to the Cloud Platform Appliance. For information, see [About Authority Delegation](#).

To add an IPv4 fixed address:

1. Navigate to the network to which you want to add a fixed address, and then select **Fixed Address** from the Add drop-down menu.  
or  
From any panel in the **DHCP** tab, expand the Toolbar and click **Add -> Fixed Address -> IPv4**.
2. In the *Add IPv4 Fixed Address* wizard, select one of the following and click **Next**:
  - **Add Fixed Address**  
or
  - **Add Fixed Address using Template**  
Click **Select Template** and select the template that you want to use. Note that when you use a template to create a fixed address, the configurations of the template apply to the new address. The appliance automatically populates the fixed address properties in the wizard. You can then edit the pre-populated properties.
    - i. Complete the following:
      1. **Network**: Click **Select Network**. When there are multiple networks, Grid Manager displays the *Select Network* dialog box from which you can select one.
      2. **IP Address**: Enter the IPv4 address for the fixed address, or click **Next Available IP** to obtain the next available IP address. For information about obtaining the next available IP address, see [About the Next Available Network or IP Address](#). Note that for Cloud Network Automation, Next Available IP is not available if the fixed address you want to create is within a delegated range.
      3. If the network of the IP address is served by a Grid member, Grid Manager displays the **Assign IP Address by** section. Select one of the following to match your criteria:



- **MAC Address:** Select this to assign a fixed address to a host with the MAC address that you specify here. Enter the MAC address in the field. For MAC address format, see [Configuring IPv4 Fixed Addresses](#).
- **DHCP Client Identifier:** Select this to assign a fixed address to a host with the DHCP client identifier that you specify here. In the field, enter the client identifier of the host to which you want the DHCP server to assign this IP address. The client identifier must be unique within the network.
  - **Match null (\0) at beginning of DHCP client identifier:** This is enabled when you select **DHCP client identifier**. Select this when a DHCP client sends a \000 prefixed to the DHCP client identifier. \0 is the null character. Some DHCP clients (for example, Microsoft) send the client identifier in a \000foo format (with the null character prefix instead of just foo). The client identifier for the requesting host and the client identifier stored in the appliance must match.
- **DHCP Relay Agent:** Select this to assign a fixed address to a host with the circuit ID or remote ID you specify here. From the drop-down list, select **Circuit ID** or **Remote ID**, and then enter the ID in the field. For information about circuit IDs and remote IDs, see [About the DHCP Relay Agent Option \(Option 82\)](#). You can enter the ID in hexadecimal format, such as ex:aa, ab, 1f:cd, or ef:23:56, or in string format, such as abcd or aa:gg. The appliance matches the value you enter here with the value sent by the DHCP client in counted octet sequence format. For information about how to use hexadecimal values, see [DHCP Option Data Types](#). The ID is case sensitive and can contain up to 230 characters. Regardless of the entry you enter here, you can define the logging format for the circuit ID and remote ID when Grid Manager displays them in the detailed lease information page. For information about how to configure the logging format, see [Defining Logging Format for DHCP Option 82](#).

Note that you cannot use the same circuit ID or remote ID for different fixed addresses if the addresses are in the same network or the same shared network.

4. **Name:** Enter a name for the Fixed Address. This field is required if the network is served by a Microsoft server. For information, see [Adding Fixed Addresses/Microsoft Reservations](#).
5. **Comment:** Optionally, enter additional information about the fixed address.
6. **Disabled:** Select this if you do not want the DHCP server to allocate this IP address at this time.

The **Cloud** section appears when the Cloud Network Automation license is installed on the Grid Master. For information, see [Deploying Cloud Network Automation](#). This section displays the following information:

- **Cloud Usage:** This field indicates whether this object is associated with any specific cloud extensible attributes or within a scope of delegation. It can be one of the following:
  - **Cloud from adapter:** Indicates that this object has been created by a cloud adapter and it may or may not be within a scope of delegation at the moment.
  - **Cloud from delegation:** Indicates that this object is within the scope of delegation or the object itself defines a scope of authority delegation, and it is not created by a cloud adapter.
  - **Used by cloud:** Indicates that this network or network container is associated with the extensible attribute **Is External** or **Is Shared** and the value is set to True, which implies the network is a private or shared network managed by the CMP, and it is not **Cloud from adapter** or **Cloud from delegation**.
  - **Non-cloud:** The object is a regular NIOS object and is not within the scope of any authority delegation nor is it associated with any of these extensible attributes: **Cloud API Owned**, **Is External** or **Is Shared**. NIOS admin users can modify this object based on their permissions.

- **Owned By:** A cloud object can be owned by the Grid Master or the cloud adapter. When the object is created by the Grid Master, this shows **Grid**. If the object is created by the cloud adapter, this shows **Adapter**.

#### Delegate authority from the Grid Master

- **Delegate To:** This field indicates whether the authority for the object you want to create has already been delegated. If so, it displays the name of the delegation.
- (Optional) Click **Next** to configure or override DHCP options as described in [About IPv4 DHCP Options](#).
  - (Applies only to Network Insight) *This step is not required for creating a new Fixed Address.* In the following Wizard step, you can optionally define the following identification values and settings for the new object's port reservation:
    - Choose the **Device Type:** **Router, Switch-Router, Switch, MSFT (Microsoft) Server, NetMRI, NIOS, VNIOS, or ESX (VMware) Server.**  
The values on this page are not required for defining the actual port reservation in a later wizard step. Certain device types could be descriptively relevant based on the type of object you are creating. As an example, the **MSFT Server** designator helps identify the new object as a Microsoft Hyper-V Host. The **ESX Server** designator can be used to identify the new object as a VMware ESX Host. These values are not required and will not affect the functionality of the object.
    - Choose the **Device Vendor:** **Cisco, Juniper, Aruba, Dell, Infoblox, or HP.**
    - You can also enter a **Location** and a **Description**. These values are advisory and not required for configuration.
    - After you define this group of settings, you can also define a device port reservation, which is done in a later step. This is not required for the Fixed Address object creation.
  - Click **Next** to initiate or disable discovery of the new Fixed Address. (*Applies only to Network Insight*) *This step is not required for creating a new Fixed Address.*
    - Choose either **Exclude from Network Discovery** or **Enable Immediate Discovery**. If you choose to Exclude, discovery will not execute on the fixed IP address. If you choose **Enable Immediate Discovery**, discovery will execute on the object after you save your settings. You may also choose to leave both options disabled.
    - By default, the new fixed address object inherits its SNMP credentials from those defined at the grid level. Should you wish to override them for a local set of credentials, check the **Override Credentials** checkbox and select the **SNMPv1/SNMPv2** or **SNMPv3** option and enter the locally used credentials.
    - You may also test the entered SNMP credentials by clicking **Test SNMP Credential**.

For descriptions of SNMP credentials for discovery, see the section [Configuring SNMP1/v2 Credentials for Polling and Configuring SNMPv3 Properties](#). These Grid-based values are inherited, by default, by each new object you create.

    - For the new object, you can check the **Override CLI Credentials** checkbox to override the inherited set of CLI credentials taken from the Grid level. This set of credentials may be used for the device that is directly associated with the new object in its port reservation.
    - You can also click **Test CLI Credentials** to enter and test a set of CLI login credentials against a device based on its IP address.  
Port control operations require CLI credentials for the involved devices. (If you are not using port control for the new object, usage of CLI credentials is optional.) Because some IPAM and DHCP objects will use port control features as part of object creation, CLI credentials are automatically leveraged as part of discovery. Ensure you have the correct sets of CLI credentials for devices in your network. For more information, see the section [Configuring CLI Discovery Properties](#).
    - SSH is the default for CLI operations. Check the **Allow Telnet** checkbox if you know the device involved in the object assignment may support Telnet but may not support SSH, or if you want Telnet as an option.

Note that all port configuration operations require CLI credentials to be entered into Grid Manager. Because some IPAM and DHCP objects will use port control features as part of object creation, CLI credentials are automatically leveraged as part of discovery. Ensure you have the correct sets of CLI credentials for devices in your network.
  - Click **Next** to define port connectivity for the device port that will be associated with the new object. *This step is not required for creating a new Fixed Address.* This feature set is also termed *portcontrol* in the NIOS/Grid

Manager system. The device whose interface the new Fixed Address will be associated should already be discovered by Network Insight.

- After choosing the device, choose the **Interface** with which the port reservation will be bound. The drop-down list shows only interfaces that are most recently found to be available by Grid Manager during the last discovery cycle.
  - The Wizard page also shows a list of any VLANs that are currently configured in the chosen device (**The following VLANs are configured**). This Wizard page allows only the assignment of an existing VLAN in the chosen device to the new port reservation.
  - Check the **Configure Port** checkbox to define port control settings for the port reservation.
  - Choose the **Data VLAN** and/or the **Voice VLAN** settings you may need for the port assignment. Depending on the selected device, you may or may not be able to apply VLAN settings.
  - Set the **Admin Status** to **Up** if you need to activate the port after assignment in the current task. All port control operations require CLI credentials to be configured. Because some IPAM and DHCP objects will use port control features as part of object creation, CLI credentials are automatically leveraged as part of discovery and definition of port configurations such as Admin Up/Down status. Ensure you have the correct sets of CLI credentials for devices in your network.
  - Enter a **Description** for the port assignment. Infoblox recommends doing so to help other technicians to recognize the port assignment task.
7. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#). When you create a new fixed address whose authority is delegated to a Cloud Platform Appliance, the required cloud extensible attributes and their values are automatically populated. You need to assign a **Subscriber Member Site** to add subscriber service related extensible attribute in order to populate Subscriber Cache.
  8. As the final step in the Add Fixed Address wizard, you define when Grid Manager creates the new object by scheduling it. You also schedule when the associated port control task executes (if a port configuration is specified).
    - To create the new object and its associated port configuration immediately, select **Now**. Grid Manager synchronizes the port reservation task to take place at the same time as the activation of the new object.
    - You can choose to have Grid Manager execute the port reservation task at the same time as the Fixed Address object creation. To do so, select **At same time as Fixed Address**.
    - You can choose to have Grid Manager execute the port reservation task at a later time by selecting **Later**. Choose a **Selected time** by entering or selecting a **Start Date** (click the calendar icon to choose a calendar date) and a **Start Time**, and choose a **Time Zone**.
  9. Choose one of the following from the **Save&...** drop-down button menu:
    - Click **Save & Close** to add the new object and close the wizard (this is the default).
    - Click **Save & Edit** to add the new object and launch the editor.
    - Click **Save & New** to add the new object and launch the wizard again to add another Fixed Address object.



#### Note

At any step during the wizard, you can click **Schedule for Later** to schedule the task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#). You cannot schedule this task when you are creating an object that is within a delegated scope. For information on viewing fixed addresses and other DHCP objects, see [Viewing IPv4 DHCP Objects](#).

## Modifying IPv4 Fixed Addresses

To modify the settings of a fixed address:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *fixed\_address* checkbox, and then click the Edit icon.
2. The *Fixed Address* editor contains the following tabs from which you can modify settings:
  - **General**: You can modify the fields, except the network address, as described in Adding IPv4 Fixed Addresses above.

- **Device Information:** Define general identity/type information for the type of device to which your new object connects. For more information, see step 5 in the previous section, [Adding IPv4 Reservations](#).
  - **Discovery:** Checking the **Enable Discovery** checkbox informs NIOS to begin discovering the fixed address after you click **Save and Close**. You manage discovery polling settings local to the fixed address from this page. For a complete overview of features on this page, see [Discovering Devices and Networks](#) and its subsections.
  - **IPv4 DHCP Options:** You can keep the inherited DHCP options or override them and enter unique settings for the fixed address. For information, see [Defining IPv4 DHCP Options](#).
  - **IPv4 Discovered Data:** Displays the discovered data of the fixed address. For information, see [Viewing Discovered Data](#).
  - **Port Reservation:** Review and edit any device port reservations that may be defined for the current object, or create a new port reservation and schedule it. For a closer look, see the section [Port Control Features in Network Insight](#), and steps 4-8 in the section [Adding IPv4 Fixed Addresses](#) above.
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#). You can edit values of inheritable extensible attributes by double clicking on the respective column. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you perform an inline editing.
  - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#).
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
    - **IPv4 DDNS:** You can keep the inherited DDNS settings or override them and enter unique settings for the fixed address. Note that you must click **Override** and select **Enable DDN Updates** for the DDNS settings you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).
    - **IPv4 BOOTP/PXE:** You can keep the inherited BOOTP properties or override them and enter unique settings for the fixed address. For information, see [Configuring IPv4 BOOTP and PXE Properties](#).
    - **Filters:** You can keep the inherited IPv4 logic filters or override them and add a new IPv4 logic filter. For information, see [Applying Filters to DHCP Objects](#).  
Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
  4. Save the configuration and click **Restart** if it appears at the top of the screen.



#### Note

At any step during the wizard, you can click **Schedule for Later** to schedule the task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#). You cannot schedule this task when you are creating an object that is within a delegated scope.

## Deleting Fixed Addresses

To delete a fixed address within the DHCP range:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *fixed\_address* checkbox or checkboxes. You cannot delete multiple fixed addresses at the same time if the authority for one of the fixed addresses is delegated to a Cloud Platform Appliance.
2. Select **Delete** from the **Delete** drop-down list.
3. In the *Delete Confirmation* dialog box, do the following:
  - **Delete associated leases with the fixed address (selected fixed IP address):** When you clear this checkbox and click **Yes**, the appliance changes the status of the associated leases from **Static** to **Active**. When you select this checkbox and click **Yes**, the appliance deletes all the leases associated with the fixed address.



#### Note

NIOS removes all the static leases associated with a fixed address when you delete a fixed address out of the DHCP range, regardless of the selection of the **Delete associated leases with the fixed address (selected fixed IP address)** checkbox in the *Delete Confirmation* dialog box.

To schedule the fixed address deletion:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *fixed\_address* checkbox.
2. Select **Schedule Delete** from the **Delete** drop-down list.
3. In the *Schedule Deletion* dialog box, complete the following:
  - **Delete Now**: Select this to delete the object upon clicking **Delete Now**.
  - **Delete Later**: Select this to schedule the deletion at a later date and time. Complete the following:
    - **Date**: Enter the date in YYYY-MM-DD (year-month-day) format. The appliance displays today's date. You can also click the calendar icon to select a date from the calendar widget.
    - **Time**: Enter the time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
    - **Time Zone**: Select a time zone for the scheduled date and time from the drop-down list. This field displays the time zone of the browser that the admin uses to log in to Grid Manager.
  - **Delete associated leases with the fixed address**: Select this checkbox to delete all the leases associated with the fixed address.
4. Click **Schedule Deletion**.

The appliance performs the deletion at the scheduled date and time, and puts all deleted objects in the Recycle Bin, if enabled. You can restore the objects if necessary.

## Configuring IPv4 Reservations

You can create a reservation as a static IP address for future use. A reservation is a pre-provisioned fixed address that is associated with a MAC address of 00:00:00:00:00:00. Since 00:00:00:00:00:00 is not a real MAC address, no client can receive this IP address from the address pool. You can reserve this static IP address and assign it to a client in the future. To create a reservation, you can do one of the following:

- Add a reservation.
- Convert a fixed address or a dynamic address with an active lease to a reservation. For information, see [Converting Objects Associated with IP Addresses](#).
- Define a fixed address with an IP address. For information, see [Adding IPv4 Fixed Addresses](#).

## Adding IPv4 Reservations



#### Note

At any step during the wizard, you can click **Schedule for Later** to schedule the task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

To create a reservation:

1. Navigate to the network to which you want to add a reservation, and then select **IPv4 Reservation** from the Add drop down menu.  
or  
From any panel in the DHCP tab, expand the Toolbar and click **Add** -> **IPv4 Reservation**.
2. In the *Add Reservation* wizard, select one of the following and click **Next**:

- **Add Reservation**  
or
  - **Add Reservation using Template**  
Click **Select Template** and select the template that you want to use. Note that when you use a template to create a reservation, the configurations of the template apply to the new address. The appliance automatically populates the reservation properties in the wizard. You can then edit the pre-populated properties.
3. Complete the following:
- **Network:** The displayed network address can either be the last selected network or the network from which you are adding the DHCP range. If no network address is displayed or if you want to specify a different network, click **Select Network**. When there are multiple networks, Grid Manager displays the *Select Network* dialog box from which you can select one.
  - **IP Address:** Enter the IP address that you want to reserve for manual assignment, or click **Next Available IP** to obtain the next available IP address. For information about obtaining the next available IP address, see [Adding IPv4 Fixed Addresses](#). Note that for Cloud Network Automation, Next Available IP is not available if the reservation you want to create is within a delegated range.
  - **Name:** Optionally, enter a name for the reservation.
  - **Comment:** Optionally, enter additional information about the reservation.
  - **Disabled:** Select this if you do not want the DHCP server to use this reservation at this time. The **Cloud** section appears when the Cloud Network Automation license is installed on the Grid Master. For information, see [Deploying Cloud Network Automation](#). This section displays the following information:
    - **Cloud Usage:** This field indicates whether this object is associated with any specific cloud extensible attributes or within a scope of delegation. It can be one of the following:
      - **Cloud from adapter:** Indicates that this object has been created by a cloud adapter and it may or may not be within a scope of delegation at the moment.
      - **Cloud from delegation:** Indicates that this object is within the scope of delegation or the object itself defines a scope of authority delegation, and it is not created by a cloud adapter.
      - **Used by cloud:** Indicates that this network or network container is associated with the extensible attribute **Is External** or **Is Shared** and the value is set to True, which implies the network is a private or shared network managed by the CMP, and it is not **Cloud from adapter** or **Cloud from delegation**.
      - **Non-cloud:** The object is a regular NIOS object and is not within the scope of any authority delegation nor is it associated with any of these extensible attributes: **Cloud API Owned**, **Is External** or **Is Shared**. NIOS admin users can modify this object based on their permissions.
      - **Owned By:** A cloud object can be owned by the Grid Master or the cloud adapter. When the object is created by the Grid Master, this shows **Grid**. If the object is created by the cloud adapter, this shows **Adapter**.
- Delegate authority from the Grid Master**
- **Delegate To:** This field indicates whether the authority for the object you want to create has already been delegated. If so, it displays the name of the delegation.
4. Click **Next** to configure or override DHCP options as described in [Defining IPv4 DHCP Options](#).
5. *(Applies only to Network Insight) This step is not required for creating a new IPv4 Reservation.* In the following Wizard step, you can optionally define the following identification values and settings for the new object's port reservation:
- Choose the **Device Type:** **Router**, **Switch-Router**, **Switch**, **MSFT (Microsoft) Server**, **NetMRI**, **NIOS**, **VNIOS**, or **ESX (VMware) Server**.  
The values on this page are not required for defining the actual port reservation in a later wizard step. Certain device types could be descriptively relevant based on the type of object you are creating. As an example, the **MSFT Server** designator helps identify the new object as a Microsoft Hyper-V Host. The **ESX Server** designator can be used to identify the new object as a VMware ESX Host. These values are not required and will not affect the functionality of the object.
  - Choose the **Device Vendor:** **Cisco**, **Juniper**, **Aruba**, **Dell**, **Infoblox**, or **HP**.
  - You can also enter a **Location** and a **Description**. These values are advisory and not required for configuration.



After you define this group of settings, you will still need to define a device port reservation, which is done in a later step.

6. *(Applies only to Network Insight)* Click **Next** to initiate or disable discovery of the new IPv4 reservation.
  - Choose either **Exclude from Network Discovery** or **Enable Immediate Discovery**. If you choose to Exclude, discovery will not execute on the object. If you choose **Enable Immediate Discovery**, discovery will execute on the object after you save your settings. You may also choose to leave both options disabled.
  - By default, the new object inherits its SNMP credentials from those defined at the grid level. Should you wish to override them for a local set of credentials, check the **Override Credentials** checkbox and select the **SNMPv1/SNMPv2** or **SNMPv3** option and enter the locally used credentials. For more information, see the sections [Configuring SNMP1/v2 Credentials for Polling](#) and [Configuring SNMPv3 Properties](#) for a complete description of SNMP credentials for discovery.
  - For the new object, you can check the **Override CLI Credentials** checkbox to override the inherited set of CLI credentials taken from the Grid level. This set of credentials may be used for the device that is directly associated with the new object (in this case, an IPv4 Reservation) in its port reservation.
  - You can also click **Test CLI Credentials** to select and test a set of CLI login credentials against a device based on its IP address.
  - Port control tasks require CLI credentials for the involved devices. (If you are not using port control for the new object, usage of CLI credentials is not required.) Because some IPAM and DHCP objects will use port control features as part of object creation, CLI credentials are automatically leveraged as part of discovery. Ensure you have the correct sets of CLI credentials for devices in your network. For more information, see the section [Configuring CLI Discovery Properties](#).
  - SSH is the default for CLI operations. Check the **Allow Telnet** checkbox if you know the device involved in the object assignment may support Telnet but may not support SSH, or if you want Telnet as an option.
7. *(Applies only with Network Insight)* Click **Next** to define optional device port association for the IPv4 reservation. *This step is optional and not required for creating the new IPv4 Reservation.* This feature set is also termed *port control* in the NIOS/Grid Manager system. The device to which the new object will be associated should already be discovered and managed from the Infoblox Grid.
  - Begin by checking the **Reserve Port** checkbox. Note that reserving a switch port does not guarantee its availability. Optionally, you can skip connecting port configuration by clicking **Next**. Click the **Clear** button to remove the selected device from the configuration.
  - Click the **Select Device** button to choose the device for which the port reservation will be associated. You should know the identity of the device to whose interface the new object will be associated before taking this step. For more information, see the section [Using the Device Selector](#).
  - After choosing the device, choose the **Interface** with which the reservation will be bound. The drop-down list shows only interfaces that are most recently found to be available by Grid Manager during the last discovery cycle.
  - The Wizard page also shows a list of any VLANs that are currently configured in the chosen device (**The following VLANs are configured**). This Wizard page allows only the assignment of an existing VLAN in the chosen device to the new port reservation.
  - Check the **Configure Port** checkbox to define specific port configuration settings for the port reservation.
  - Choose the **DataVLAN** and/or the **VoiceVLAN** settings you may need for the port assignment. Depending on the selected device, you may or may not be able to apply VLAN settings.
  - Set the **AdminStatus** to **Up** if you need to activate the port after assignment in the current task. All port control operations require CLI credentials to be entered into Grid Manager. Because some IPAM and DHCP objects will use port control features as part of object creation, CLI credentials are automatically leveraged as part of discovery and definition of port configurations such as Admin Up/Down status. Ensure you have the correct sets of CLI credentials for devices in your network.
  - Enter a **Description** for the port reservation. Infoblox recommends doing so to help other technicians to recognize the port assignment task.
8. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#). You need to assign a **Subscriber Member Site** to add subscriber service related extensible attribute in order to populate Subscriber Cache.
9. As the final step in the Add IPv4 Reservation wizard, you define when Grid Manager creates the new object by scheduling it. You also schedule when the associated port control task executes (if a port configuration is specified).



- To create the new IPv4 Reservation and its associated port reservation immediately, select **Now**. Grid Manager synchronizes the port control task to take place at the same time as the activation of the new object.
  - You can choose to have Grid Manager execute the port control task at the same time as the object creation. To do so, select **At same time as IPv4 Reservation**.
  - You can choose to have Grid Manager execute the port control task at a later time by selecting **Later**. Choose a **Selected time** by entering or selecting a **Start Date** (click the calendar icon to choose a calendar date) and a **Start Time**, and choose a **Time Zone**.
10. Choose one of the following from the **Save&...** drop-down button menu:
    - Click **Save & Close** to add the new object and close the wizard (this is the default).
    - Click **Save & Edit** to add the new object and launch the editor.
    - Click **Save & New** to add the new object and launch the wizard again to add another IPv4 Reservation object.
  11. Click **Restart** if it appears at the top of the screen.



#### Note

At any step during the wizard, you can click **Schedule for Later** to schedule the task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#). You cannot schedule this task when you are creating an object that is within a delegated scope.

## Modifying IPv4 Reservations

To modify an IPv4 reservation:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *reservation* checkbox, and then click the Edit icon.
2. The *Reservation Address* editor contains the following tabs from which you can modify data:
  - **General**: Modify the fields, except the network address, as described in Adding IPv4 Reservations above.
  - **Device Information**: Define general identity/type information for the type of device to which your new object connects. For more information, see Step 5 in the previous section.
  - **IPv4 DHCP Options**: Keep the inherited DHCP options or override them and enter unique settings for the reservation. For information, see [Defining IPv4 DHCP Options](#).
  - **IPv4 Discovered Data**: Displays the discovered data of the reservation. For information, see [Viewing Discovered Data](#).
  - **Port Reservation**: Review and edit any device port reservations that may be defined for the current object, or create a new port reservation and schedule it. For a closer look, see the section [Port Control Features in Network Insight](#), and steps 4-8 in the section Adding IPv4 Reservations above.
  - **Extensible Attributes**: Add and delete extensible attributes that are associated with a reservation. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#). You can edit values of inheritable extensible attributes by double clicking on the respective column. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you perform an inline editing. The *Descendant Actions* dialog box is displayed when you click **Save**. For information, see [Managing Inheritable Extensible Attributes at the Parent and Descendant Level](#). If you delete the value of an inheritable extensible attribute at the parent level, you can choose to preserve the descendant value or remove it. For information, see [Deleting Inheritable Extensible Attributes Associated with Parent Objects](#).
  - **Permissions**: This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#).
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
  - **IPv4 DDNS**: Keep the inherited DDNS settings or override them and enter unique settings for the reservation. Note that you must click **Override** and select **Enable DDN Supdates** for the DDNS settings

you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).

- **IPv4 BOOTP/PXE:** You can keep the inherited BOOTP properties or override them and enter unique settings for the reservation. For information, see [Configuring IPv4 BOOTP and PXE Properties](#).
- **Filters:** You can keep the inherited IPv4 logic filters or override them and add a new IPv4 logic filter. For information, see [Applying Filters to DHCP Objects](#).

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.

4. Save the configuration and click **Restart** if it appears at the top of the screen.



#### Note

At any step during the wizard, you can click **Schedule for Later** to schedule the task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#). You cannot schedule this task when you are creating an object that is within a delegated scope.

## Viewing IPv4 DHCP Objects

To view the address ranges, fixed addresses and reservations in a network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr\_range*.
2. Grid Manager displays the following information:
  - **IP Address:** The IP address of the object in the DHCP range. For exclusion ranges, this displays the start and end IP addresses. For host records with multiple IP addresses, each IP address is displayed separately. The appliance highlights disabled DHCP objects in gray. A DHCP object can be a fixed address, reservation, host configured for DHCP, or roaming host with an allocated IP address.
  - **Type:** The object type, such as **Fixed Address**.
  - **Name:** The object name. For example, if the IP address belongs to a host record, this field displays the hostname.
  - **Fingerprint:** The name of the DHCP fingerprint or vendor ID of the network device that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#).
  - **Comment:** The information you entered for the object.
  - **Site:** The site to which the object belongs. This is one of the predefined extensible attributes. You can edit values of inheritable extensible attributes by double clicking on the respective column. If an extensible attribute has an inherited value, then the cell is highlighted in blue when you perform an inline editing.

You can select **Disabled** or available extensible attributes for display.

You can also do the following:

- Sort the data in ascending or descending order by column.
- Create a bookmark for the range.
- Delete or schedule the deletion of a selected object or multiple objects in the range.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Select an object and view detailed information.
- Print or export the data.

## About Roaming Hosts

A roaming host is a host with a dynamically assigned IP address and a specific set of properties and DHCP options. When you create a roaming host for a network device, the device can receive any dynamically assigned address from the network to which it belongs. You can create roaming hosts for devices, such as laptop computers and mobile phones, that require different IP addresses each time they are moved from one network to another and require a unique set of DHCP options.

You can configure IPv4 addresses, IPv6 addresses, or IPv4 and IPv6 addresses for roaming hosts that require both types of addresses. When you configure IPv4 addresses for a roaming host, you must specify the host MAC address or a DHCP client identifier that the appliance uses to match the host, and specify DHCP options for the host. The appliance assigns an IP address from the DHCP range associated with the network from which the address request originates. You can configure an IPv6 prefix or address for a DHCP client. When you do, you must specify the DUID of the host so the appliance can use the DUID to match the host.

A roaming host also receives DHCP options from the Grid, member, network, or shared network with which it associates. When you configure a roaming host, you must configure it in a specific network view. If you have multiple network views, you must specify the network view to which the requesting hosts belong so the appliance can assign addresses to the hosts from the networks within the same network view.

After you enable support for roaming hosts at the Grid level, you can add a roaming host that supports IPv4, IPv6, or both protocols. You can also convert an IPv4 roaming host to an IPv6 roaming host and vice versa, or convert an IPv4 or IPv6 roaming host to one that supports both IPv4 and IPv6.

## Configuring Roaming Hosts

To configure a roaming host, perform the following tasks:

1. Enable support for roaming hosts at the Grid level. For information, see, [Enabling Support for Roaming Hosts below](#).
2. Add a roaming host.
  - To add an IPv4 roaming host, see [Adding IPv4 Roaming Hosts below](#)
  - To add an IPv6 roaming host, see [Adding IPv6 Roaming Hosts below](#)
  - To add a dual stack roaming host, see [Adding IPv4/IPv6 Roaming Hosts below](#)
  - Optionally, configure DHCP properties for the roaming host. You can override properties set for the upper levels and enter unique values for the roaming hosts. For information, see [Defining IPv4 DHCP Options](#).

You can do the following after you configure roaming hosts:

- View the configured roaming hosts. For information, see [Viewing Roaming Hosts below](#).
- Modify existing roaming hosts. For information, see [Setting Properties for Roaming Hosts below](#).
- Delete roaming hosts that are not currently in use. For information, see [Deleting Roaming Hosts below](#).

## Enabling Support for Roaming Hosts

You must first enable support for roaming hosts before adding them. After you enable this feature, you can disable it only after you delete all the existing roaming hosts.

To enable support for roaming hosts:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Expand the Toolbar and click **Grid DHCP Properties**.
3. In the *General Advanced* tab, select **Enable support for roaming host**.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Adding IPv4 Roaming Hosts

To add an IPv4 roaming host:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Select a network view from the drop-down list.
3. Expand the Toolbar and click **Add -> Roaming Host -> IPv4**.
4. In the *Add Roaming Host* wizard, select one of the following and click **Next**:
  - **Add Roaming Host**  
or
  - **Add Roaming Host using Template**  
Click **Select Template** to create a roaming host using a fixed address/reservation template. In the *DHCP Template Selector* dialog box, select the template that you want to use. Note that when you use a template to create a roaming host, the configurations of the template apply to the new host. The appliance automatically populates the host properties in the wizard. You can then edit the pre-populated properties.

5. Complete the following:
  - **Name:** Enter the name of the roaming host. The name must be unique for each roaming host in a given network view.
  - **Assign IPv4 Address by:** Select one of the following criteria on which the appliance matches when assigning an IP address to the host.
    - **MAC Address:** Select this to assign a dynamic IP address to a host, provided that the MAC address of the requesting host matches the MAC address that you specify here.
    - **DHCP Client Identifier:** Select this to assign a dynamic IP address to a host with the same DHCP client identifier that you specify here. When you select this, the **Match null (\0) at beginning of DHCP client identifier** checkbox is displayed. Select this when a DHCP client sends a \000 prefixed to the DHCP client identifier. \0 is the null character. Some DHCP clients (for example, Microsoft) send the client identifier in a \000foo format (with the null character prefix instead of just foo). The client identifier for the requesting host and the client identifier stored in the appliance must match.
  - **Comment:** Enter useful information about the roaming host.
  - **Disabled:** Select this if you do not want the DHCP server to use this roaming host definition. When you disable a roaming host, the host gets an IP address without the defined DHCP options.
6. Click **Next** to configure the DHCP options for the roaming host, as described in [Defining IPv4 DHCP Options](#).
7. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).
8. Save the configuration and click **Restart** if it appears at the top of the screen.  
or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *ScheduleChange* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Adding IPv6 Roaming Hosts

To add an IPv6 roaming host:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Select a network view from the drop-down list.
3. Expand the Toolbar and click **Add -> Roaming Host -> IPv6**.
4. In the *Add Roaming Host* wizard, select one of the following and click **Next**:
  - **Add IPv6 Roaming Host**
  - or
  - **Add Roaming Host Using IPv6 Template**

Click **SelectIPv6Template** to create a roaming host using an IPv6 fixed address template. In the *DHCPTemplateSelector* dialog box, select the template that you want to use. Note that when you use a template to create a roaming host, the configurations of the template apply to the new host. The appliance automatically populates the host properties in the wizard. You can then edit the pre-populated properties.

5. Complete the following:
  - **Name:** Enter the name of the roaming host. The name must be unique for each roaming host in a given network view.
  - **DUID:** Enter the DHCP unique identifier of the host.
  - **Comment:** Optionally, enter additional information about the roaming host.
  - **Disabled:** Select this if you do not want the DHCP server to use this roaming host definition. When you disable a roaming host, the host gets an IP address without the defined DHCP options.
6. Click **Next** to configure the DHCP options for the roaming host, as described in [Defining General IPv6 Properties](#).
7. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).
8. Save the configuration and click **Restart** if it appears at the top of the screen.  
or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Adding IPv4/IPv6 Roaming Hosts

To add an IPv4/IPv6 roaming host:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Select a network view from the drop-down list.
3. Expand the Toolbar and click **Add -> Roaming Host -> Both**.
4. In the *Add Roaming Host* wizard, select one of the following and click **Next**:
  - **Add Roaming Host**  
or
  - **Add Roaming Host using Both IPv4 and IPv6 Templates**  
When you use both templates to create a roaming host, the appliance applies the IPv4 template and then the IPv6 template. Therefore, the comments and extensible attributes from the IPv6 template override those from the IPv4 template.
5. Complete the following:
  - **Name**: Enter the name of the roaming host. The name must be unique for each roaming host in a given network view.
  - **Assign IP Address by**: Select one of the following criteria on which the appliance matches when assigning an IP address to the host.
    - **MAC Address**: Select this to assign a dynamic IP address to a host, provided that the MAC address of the requesting host matches the MAC address that you specify here.
    - **DHCP Client Identifier**: Select this to assign a dynamic IP address to a host with the same DHCP client identifier that you specify here. When you select this, the **Match null (0) at beginning of DHCP client identifier** checkbox is displayed. Select this when a DHCP client sends a \000 prefixed to the DHCP client identifier. \0 is the null character. Some DHCP clients (for example, Microsoft) send the client identifier in a \000foo format (with the null character prefix instead of just foo). The client identifier for the requesting host and the client identifier stored in the appliance must match.
    - **DUID**: Specify the DHCP unique identifier of the host.
    - **Comment**: If both IPv4 and IPv6 templates were used to create the host, this field displays the comment from the IPv6 template. You can change or add information.
    - **Disabled**: Select this if you do not want the DHCP server to use this roaming host definition. When you disable a roaming host, the host gets an IP address without the defined DHCP options.
6. Click **Next** to configure the IPv4 DHCP options for the roaming host, as described in [Defining IPv4 DHCP Options](#).
7. Click **Next** to configure IPv6 properties described in [Defining General IPv6 Properties](#).
8. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. If both IPv4 and IPv6 templates were used to create the host, this panel displays the attributes from the IPv6 template. You can change or add information. For information, see [Using Extensible Attributes](#).
9. Save the configuration and click **Restart** if it appears at the top of the screen.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *ScheduleChange* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Viewing Roaming Hosts

To view a list of roaming hosts in a specific network view:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Roaming Hosts**.
2. From the Network View drop-down list, select the network view to which the roaming hosts belong.
3. The Grid Manager displays the following for each roaming host:
  - **Name**: The name of the roaming host.
  - **Address**: The IP address of the roaming host.
  - **Comment**: The information that you entered for the roaming host.
  - **Site**: The site to which the template belongs. This is one of the predefined extensible attributes.

You can select **Disabled** and available extensible attributes for display.

You can also do the following:

- Sort the displayed data in ascending or descending order by column.
- Use filters and the **Goto** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Goto** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).

## Setting Properties for Roaming Hosts

You can modify an existing roaming host to add, modify or delete IPv4 or IPv6 addresses, and to set IPv4 and IPv6 DHCP properties.

1. From the **DataManagement** tab, select the **DHCP** tab -> **Networks** tab -> **RoamingHosts** section -> *roaming\_host* checkbox, and then click the Edit icon.
2. The *RoamingHost* editor contains the following tabs from which you can modify data:
  - **General**: Edit the fields as described in Adding IPv4 Roaming Hosts above, except for the **Templates** field.
  - **IPv4 DHCP Options**: Keep the inherited DHCP options or override them and enter unique settings for the roaming host. For information, see [Defining IPv4 DHCP Options](#).
  - **IPv6 DHCP Options**: Keep the inherited IPv6 DHCP properties or override them. For more information, see [Defining General IPv6 Properties](#).
  - **Extensible Attributes**: Add and delete extensible attributes that are associated with a roaming host. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions**: This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#).
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
  - **IPv4DDNS**: Click **Override** and select **Enable DDNS updates** for the DDNS settings you configure in this tab to take effect. You can specify the following:
    - **DDNS Domain Name**: Specify the domain name that the appliance uses to update DNS.
    - **DDNS Hostname**: Select the Replace the host name dynamically provided by the client/member with the roaming host name checkbox to use the name of the roaming host record as the name of the client for DDNS updates.

For information about DDNS, see [Configuring DDNS Updates](#).

- **IPv4 BOOTP/PXE**: Keep the inherited PXE and BOOTP properties or override them and enter unique settings for the roaming host. For information, see [Configuring DHCP for IPv4](#).
- **IPv6 DDNS**: Click **Override** and select **Enable DDNS Updates** for the DDNS settings you configure in this tab to take effect. You can specify the following:
  - **DDNS Domain Name**: Specify the domain name that the appliance uses to update DNS.
  - **DDNS Hostname**: Select the Replace the host name dynamically provided by the client/member with the roaming host name checkbox to use the name of the roaming host record as the name of the client for DDNS updates.

For information about DDNS, see [Configuring DDNS Updates](#).

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

You can also click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).



## Deleting Roaming Hosts

To delete a roaming host:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Roaming Hosts** -> *roaming\_host* checkbox, and then select **Delete** or **Schedule Delete** from the drop-down menu.
2. To delete the roaming host now, in the *Delete Confirmation* dialog box, click **Yes**. To schedule the deletion, see [Scheduling Deletions](#). The Grid Manager puts the deleted roaming host in the Recycle Bin, if enabled.

## Managing DHCP Templates

A template contains a set of predefined properties that you use to create IPv4 and IPv6 DHCP objects. It is metadata that you can modify and reuse. Using a template enables you to create objects in a quick and consistent way. You can define the object properties once in a template, and then create multiple objects that inherit their properties from the template. For example, you can create a network template that has a fixed netmask of /24 and extensible attribute "State" set to California. You can then use the template to create networks in California that contain /24 netmasks. You can also modify and delete a template. Note that modifying or deleting a template does not affect existing objects created based on the template. You must be a superuser or have read/write permissions to add, modify, or delete a template. A superuser can set other admin group privileges on templates. For information, see [Administrative Permissions for IPv4 or IPv6 DHCP Templates](#). You can also define extensible attributes for these templates when you create them. For information, see [Using Extensible Attributes](#).

This section explains how to configure and manage IPv4 and IPv6 DHCP templates. It contains the following topics:

- [About IPv4 DHCP Templates](#)
- [About IPv6 DHCP Templates](#)
- [About IPv6 Range Templates](#)
- [About IPv6 Fixed Address Templates](#)
- [About IPv6 Network Templates](#)
- [Viewing and Deleting Templates](#)

## About IPv4 DHCP Templates

You can use templates to create DHCP IPv4 ranges, fixed addresses, reservations, roaming hosts, and networks. You can create the following IPv4 templates:

- A DHCP range template, containing DHCP range settings, such as the total number of IP addresses allocated to a range. You can add a DHCP range template to a network template. For information, see [About IPv4 Range Templates](#) below.
- A fixed address/reservation template, containing information for creating fixed addresses, reservations, or roaming hosts. You can add a fixed address/reservation template to a network template. For information, see [About IPv4 Network Templates](#) below.
- A network template, containing basic network properties for creating networks. It is also a container that holds your DHCP range templates and fixed address/reservation templates. When you create a network using a network template, the network inherits the properties of the range and fixed address/reservation templates. You can create a network in any network view using a network template. For information, see [About IPv4 Network Templates](#) below.

Because you can potentially add DHCP range and fixed address/reservation templates to a network template, create the DHCP range and fixed address/reservation templates before you create a network template. For information, see [Configuration Example: Creating an IPv4 Network Using a Template](#) below.

## About IPv4 Range Templates

When you create an IPv4 range template, the start and end address fields are based on the specified offset from the network start address and the number of IP addresses in the range. After you create a DHCP range template, you can



configure additional properties such as exclusion ranges and DHCP filters, as described in [Modifying IPv4 Range Templates](#) below. Then when you use the template to create a DHCP range, the range inherits the properties of the template. You can also include a DHCP range template in a network template to automatically create a DHCP range when you use that network template.

If you have deployed the Cloud Network Automation license on the Grid Master, you can configure range templates for cloud delegation. When you select a default Cloud Platform Appliance for a template, all ranges you create using this template will delegate authority to the same Cloud member. Note that when a Cloud member is removed from the Grid, the delegation will also be removed from the template. For information about Cloud Network Automation, see [Deploying Cloud Network Automation](#).

## Adding IPv4 Range Templates

To create an IPv4 DHCP range template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab, and then expand the Toolbar and click **Add -> Templates -> Range -> IPv4**.
2. In the *Add IPv4 Range Template* wizard, do the following:
  - **Name:** Enter a name that helps identify the DHCP range template. For example, enter **Region 1 IT** if you want to use this template to create DHCP ranges for the IT department in Region 1.
  - **Offset:** An offset in a DHCP range template determines the starting IP address of the range. The appliance adds the offset value you enter here to the start IP address of the network in which you create a DHCP range using this template. That IP address becomes the start IP address of the DHCP range. For example, you specify an offset value of 25 for a 25.0.0.0/8 network using the DHCP range template, the appliance creates a DHCP range with the start IP address of 25.0.0.25 in the network.
  - **Number of Addresses:** Enter the total number of IP addresses to be included in the DHCP range.
  - **Comment:** Enter useful information about the template.

The **Cloud** section appears when the Cloud Network Automation license is installed on the Grid Master. For information, see [Deploying Cloud Network Automation](#). To configure this template for cloud delegation, complete the following:

**Use for cloud delegation:** Select this checkbox to enable cloud delegation for this template.

### Delegate authority from the Grid Master

- **Delegate To:** In a non-cloud API request, this parameter defines the default member to which authority is delegated. In a cloud API request, the appliance ignores this parameter, which allows you to use this template to create an object on different Cloud Platform Appliances. Click **Select** to choose the default Cloud Platform Appliance to which you want to delegate authority. The *Member Selector* displays only Cloud Platform Appliances in the Grid. Click the member, and Grid Manager displays the member name next to this field.
3. Click **Next** and select one of the following to provide DHCP services for the range:
    - **None (Reserved Range):** Select this if you want to reserve this address range for static hosts. Addresses in this range cannot be allocated as dynamic addresses. You can allocate the next available IP from this range to a static host. This is selected by default.
    - **Grid Member:** Click **Select** and choose a Grid member from the drop-down list.
    - **Failover Association:** Click **Select** and choose a failover association. Only failover associations that provide DHCP services in the network view of the DHCP range appear in the drop-down list.
    - **Microsoft DHCP Server:** Click **Select** and choose a Microsoft server from the drop-down list. The drop-down list displays only the servers that are associated with the network to which the DHCP range belongs.
  4. Click **Next** to configure or override DHCP options as described in [Defining IPv4 DHCP Options](#).
  5. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).
  6. Save the configuration and click **Restart** if it appears at the top of the screen.

## Modifying IPv4 Range Templates

After you use the wizard to create an IPv4 DHCP range template, you can set additional properties for the template. Following are some guidelines:

- In the **DHCP Options** tab of a DHCP range template, the broadcast address is an address offset number rather than a broadcast IP address; network router addresses are offset numbers as well. An offset in a DHCP range template indicates the starting IP address of the DHCP range object created from the template. For example, you can create a network template called *test\_network\_template* and a DHCP range template *test\_range\_template* linked to this network template. If the *test\_range\_template* has an offset value 10, when you create a 10.0.0.0/8 network using the *test\_network\_template*, the appliance creates a DHCP range with the starting IP address 10.0.0.10. If you create a 20.0.0.0/8 network using the *test\_network\_template*, the appliance creates a DHCP range with the starting IP address 20.0.0.10

For the exclusion range in the template, the start and end addresses are determined by the number of offsets in the DHCP range template's start address and the number of IP addresses in the exclusion range. For more information about exclusion ranges, see [About DHCP Ranges](#).

To modify and set properties for a DHCP range template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* checkbox, and then click the Edit icon.
2. The *DHCP Range Template* editor contains the following tabs from which you can modify data:
  - **General**: Modify general information described in the previous section Adding IPv4 Range Templates.
  - **Member Assignment**: Change the Grid member, failover association, or Microsoft server that provides DHCP services for this template. You can also add or delete a member or failover association. For information, see [Adding IPv4 Address Ranges](#).

When you change the member assignment of DHCP ranges from **Failover Association** to **Grid Member** and then back to **Failover Association**, the leases in the primary and secondary servers fall out of sync. To resynchronize the peers, the failover association of the secondary server is put in the Recover-Wait state for the entire duration of Maximum Client Lead Time while a forced recovery takes place. During this period, only the IP addresses allocated to the primary server are available for lease.

- **IPv4 DHCP Options**: Keep the inherited DHCP options or override them and enter unique settings for the template. For information, see [Defining IPv4 DHCP Options](#).
  - **Extensible Attributes**: Add and delete extensible attributes that are associated with this template. You can also modify the values of the extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions**: This tab appears only if you belong to a superuser admin group. For information, see [About Administrative Permissions](#).
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify data:
    - **IPv4DDNS**: Keep the inherited DDNS settings or override them and enter unique settings for the template. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).
    - **IPv4 BOOTP/PXE**: Keep the inherited BOOTP properties or override them and enter unique settings for the template. For information, see [Configuring IPv4 BOOTP and PXE Properties](#).
    - **Exclusion Ranges**: Configure a range of IP addresses that the appliance does not use for dynamic address assignments. Complete the following:
      - **Offset**: An offset for an exclusion range determines the start IP address of the exclusion range. The appliance adds the offset value you enter here to the start IP address of the DHCP range created using this template. That IP address becomes the start IP address of the exclusion range.
      - **Number of Addresses**: Enter the number of IP addresses to be included in the exclusion range.
      - **Comment**: Enter useful information about the exclusion range.
    - **IPv4 DHCP Thresholds**: Keep the inherited thresholds settings or override them and enter unique settings for the template. For information, see [Configuring Thresholds for DHCP Ranges](#).
  4. Save the configuration and click **Restart** if it appears at the top of the screen.

 Note

Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.

## About IPv4 Fixed Address/Reservation Templates

You can use an IPv4 fixed address/reservation template to create fixed addresses, reservations and roaming hosts. When you create an IPv4 fixed address/reservation template, you can specify an offset and number of addresses. This is used when you include the template in a network template. When you include a fixed address/reservation template in a network template, the DHCP server automatically creates reservations based on the offset and number of addresses you specified in the fixed/address reservation template. It does not create fixed addresses.

After you create a fixed address/reservation template using the wizard, you can configure additional properties as described in [Modifying IPv4 Fixed Address/Reservation Templates](#) below. Then when you use the template to create a fixed address, it inherits the properties of the template.

If you have deployed the Cloud Network Automation license on the Grid Master, you can configure fixed address templates for cloud delegation. When you configure a template for cloud delegation, all fixed addresses you create using this template will inherit authority delegations from their parent objects. For information about Cloud Network Automation, see [Deploying Cloud Network Automation](#).

## Adding IPv4 Fixed Address/Reservation Templates

To create an IPv4 fixed address/reservation template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** section.
2. Click the Add drop-down list and select **IPv4 Fixed Address/Reservation Template**.
3. In the *Add IPv4 Fixed Address/Reservation Template* wizard, enter the following:
  - **Name:** Enter a name that helps identify the fixed address/reservation template. For example, you can enter **HP Printer** when you create a template that contains settings for assigning fixed addresses or reservations to HP printers.
  - **Comment:** Optionally, enter additional information about the template.
  - **Use for cloud delegation:** When you select this checkbox, all fixed addresses you create using this template inherit authority delegation from their parent objects.

In the **Optional Settings For Range of Objects** section, do the following:

- **Offset:** An offset in a fixed address/reservation template determines the start IP address of the object created from the template. The appliance adds the offset value you enter here to the start IP address of the network in which you create objects using this template. That IP address becomes the start IP address of the object.
  - **Number of Addresses:** Enter the number of IP addresses to be used as fixed addresses, reservations, or roaming hosts.  
Note that the appliance uses the offset and number of addresses only when this template is used in a network template.
4. Click **Next** to configure or override DHCP options as described in [Defining IPv4 DHCP Options](#).
  5. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).
  6. Save the configuration and click **Restart** if it appears at the top of the screen.

## Modifying IPv4 Fixed Address/Reservation Templates

To modify a fixed address/reservation template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* checkbox, and then click the Edit icon.
2. The *Fixed Address/Reservation Template* editor contains the following tabs from which you can modify data:
  - **General:** Modify general information for the template as described in [Adding IPv4 Fixed Address/Reservation Templates](#) above.
  - **IPv4 DHCP Options:** Keep the inherited DHCP options or override them and enter unique settings for the template. For information, see [Defining IPv4 DHCP Options](#).

- **Extensible Attributes:** Add and delete extensible attributes that are associated with the template. You can also modify the values of the extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [About Administrative Permissions](#).
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify data:
    - **IPv4 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the template. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).
    - **IPv4 Filters:** You can keep the inherited IPv4 logic filters or override them and add a new IPv4 logic filter for the template. For information, see [Applying Filters to DHCP Objects](#).
    - **IPv4 BOOTP/PXE:** Keep the inherited BOOTP properties or override them and enter unique settings for the template. For information, see [Configuring IPv4 BOOTP and PXE Properties](#).  
Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
  4. Save the configuration and click **Restart** if it appears at the top of the screen.

## About IPv4 Network Templates

You can create IPv4 network templates to facilitate network configuration. You can use network templates to create networks in any network view. When you create a network template, you do not specify a network address. You enter the network address when you create an actual network from the template. You can specify a netmask or allow the user to define the netmask when they create the actual network.

A network template is useful for setting up a network with fixed addresses and DHCP ranges already defined. You can add DHCP range or fixed address/reservation templates to a network template. Once the fixed address and DHCP range information is set up, the network template contains a range template list and a fixed address/reservation template list. When you enable support for RIR updates, you can create IPv4 network templates specific for RIR associated networks. For information about RIR updates, see [RIR Registration Updates](#).

If you have deployed the Cloud Network Automation on the Grid Master, you can configure network templates for cloud delegation. When you select a default Cloud Platform Appliance for a template, all networks you create using this template will delegate authority to the same Cloud member. If you want to associate any range or fixed address templates with a network template, ensure that you enable "**User for cloud delegation**" for the network, range, and fixed address templates. Note that when a Cloud member is removed from the Grid, the delegation will also be removed from the template. For information about Cloud Network Automation, see [Deploying Cloud Network Automation](#).

## Adding IPv4 Network Templates

To create a network template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** section.
2. Click the drop-down menu of the Add icon and select **IPv4 Network Template**.
3. In the *Add IPv4 Network Template* wizard, do the following:
  - **Regional Internet Registry:** This section appears only when support for RIR updates is enabled. For information about RIR, see [RIR Registration Updates](#). Complete the following to create a network template for an RIR IPv4 network container or network:
    - **Internet Registry:** Select the RIR from the drop-down list. The default is **RIPE**. When you select **None**, the network is not associated with an RIR organization.
    - **Organization ID:** Click **Select Organization** and select an organization from the *RIROrganizationSelector* dialog box.
    - **Registration Status:** The default is **Not Registered**. When using this template to add an RIR allocated network, you can change this to **Registered** and select the **Do not update registrations** checkbox below. Note that when you select **API** as the communication method, the registration status will be updated automatically after the registration update is completed. However, when you select **Email** as the communication method, the registration status will not be automatically updated. If you are creating a new network and the registration update is completed successfully, the status will be changed to **Registered**. If the update fails, the status will be changed to **Not Registered**.
    - **Registration Action:** Select the registration action from the drop-down list. When you select **Create**, the appliance creates the IPv4 network and assigns it to the selected organization. When you select **None**, the appliance does not send registration updates to RIPE. When you use this template to add an existing RIR allocated network to NIOS, select **None**. When you use this

template to add networks to an RIR allocated network (a parent network), select **Create**. Ensure that the parent network associated with an RIR organization already exists.

- **Do not update registrations:** Select this checkbox if you do not want the appliance to submit RIR updates to RIPE. By default, the appliance sends updates to the RIR database based on the configured communication method.
- **Name:** Enter a name that helps identify the network template. For example, you can enter **Class C** if you want to configure the template for creating Class C networks.
- **Netmask:** Select one of the following options:
  - **Fixed:** Select this and adjust the netmask slider to a fixed netmask for this network template. When you select this option, users cannot specify another netmask when they use this template to create a network. For example, if you select /24 as the fixed netmask, all networks created using this template have a /24 netmask.
  - **Allow User to Specify Netmask:** Select this to allow users to specify the subnet mask when creating networks using this template.
- **Comment:** Optionally, enter additional information about the template.
- **Automatically Create Reverse-Mapping Zone:** This function is enabled if the fixed netmask of the template equals /8, /16, and /24, or if you select the **Allow User to Specify Netmask** option. Select this if you want the appliance to automatically create the corresponding reverse-mapping zone for the networks created using this template. A reverse-mapping zone is an area of network space for which one or more name servers have the responsibility for responding to address-to-name queries. These zones are created in the DNS view assigned to receive dynamic DNS updates at the network level.

The **Cloud** section appears when the Cloud Network Automation license is installed on the Grid Master. For information, see [Deploying Cloud Network Automation](#). To configure this template for cloud delegation, complete the following:

**Use for cloud delegation:** Select this checkbox to enable cloud delegation for this template.

When you select this for the network template, all range and fixed address templates that you want to associate with this network template must also be enabled for "**Use for cloud delegation.**"

#### Delegate authority from the Grid Master

**Delegate To:** In a non-cloud API request, this parameter defines the default member to which authority is delegated. In a cloud API request, the appliance ignores this parameter, which allows you to use this template to create an object on different Cloud Platform Appliances. Click **Select** to choose the default Cloud Platform Appliance to which you want to delegate authority. The *Member Selector* displays only Cloud Platform Appliances in the Grid. Click the member, and Grid Manager displays the member name next to this field.

4. Click **Next** and do the following to assign either Grid members or Microsoft DHCP servers to this network template. Ensure that you include members or Microsoft servers that are associated with other templates that you plan to add to this network template. You can assign one or multiple members to this template. However, you cannot assign a combination of NIOS Grid members and vNIOS Grid members to the template. You can also assign multiple Microsoft servers to a template, but you cannot assign a mix of Microsoft servers and Grid members to a template.
  - click the Add icon and select one of the following options:
    - **Add Infoblox Member:** Select this option to add a Grid member as a DHCP server for the networks created using this template. Select the Grid member from the *Member Selector* dialog box. Keep in mind, DHCP properties for the network are inherited from this member. Networks created using this template can be served by multiple members, but a member can serve networks in one network view only.  
or
    - **Add Microsoft Server:** Select this option to add a Microsoft server as a DHCP server for the networks created using this template. Select the Microsoft server from the *Microsoft Server Selector* dialog box.
5. Click **Next** to associate Active Directory Sites with the network. For more information, see [Associating Active Directory Sites with Networks](#).
6. Click **Next** and do the following to include IPv4 address range and fixed address/reservation templates in the network template. Note that when you select a fixed address/reservation template, only reservations, not fixed



addresses, are created for networks created using this template. You cannot add a fixed address/reservation template that does not contain an offset value or a total number of IP addresses for a range.

- a. Click the Add icon.
  - b. In the *DHCP Template Selector* dialog box, choose the template that you want to include in this network template. You can choose a DHCP range or fixed address/reservation template. Use SHIFT+click and CTRL+click to select multiple templates.
  - c. Click the Select icon.  
You can delete a template from the table by selecting it and clicking the Delete icon.
7. Click **Next** to configure or override DHCP options as described in [Defining IPv4 DHCP Options](#).
  8. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).  
If you are adding an RIR network, the RIR network attribute table appears. For information about these attributes and how to enter them, see [RIR Network Attributes](#).
  9. Save the configuration and click **Restart** if it appears at the top of the screen.

## Modifying IPv4 Network Templates

To modify a network template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* checkbox, and then click the Edit icon.
2. The *IPv4 Network Template* editor contains the following tabs from which you can modify data:
  - **General**: Modify general information described in Adding IPv4 Network Templates above.
  - **Member Assignment**: Change the Microsoft servers or Grid members that provide DHCP services for this template. For information, see [Adding IPv4 Networks](#).
  - **Templates**: Add or delete DHCP range and fixed address/reservation templates. For information, see About IPv4 Range Templates and About IPv4 Fixed Address/Reservation Templates above.
  - **IPv4 DHCP Options**: Keep the inherited DHCP options or override them and enter unique settings for the template. For information, see [Defining IPv4 DHCP Options](#).
  - **RIR Registration**: Modify RIR network information. This tab appears only when support for RIR updates is enabled. For information, see [Modifying RIR Network Data](#).
  - **Extensible Attributes**: Add and delete extensible attributes that are associated with the template. You can also modify the values of the extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions**: This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#).
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify data:
  - **IPv4DDNS**: Keep the inherited DDNS settings or override them and enter unique settings for the template. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).
  - **IPv4 BOOTP/PXE**: Keep the inherited BOOTP properties or override them and enter unique settings for the template. For information, see [Configuring IPv4 BOOTP and PXE Properties](#).
  - **Filters**: You can keep the inherited IPv4 logic filters or override them and add a new IPv4 logic filter for the template. For information, see [Applying Filters to DHCP Objects](#).
  - **IPv4 DHCP Thresholds**: Keep the inherited thresholds settings or override them and enter unique settings for the template. For information, see [Configuring Thresholds for DHCP Ranges](#).  
Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuration Example: Creating an IPv4 Network Using a Template

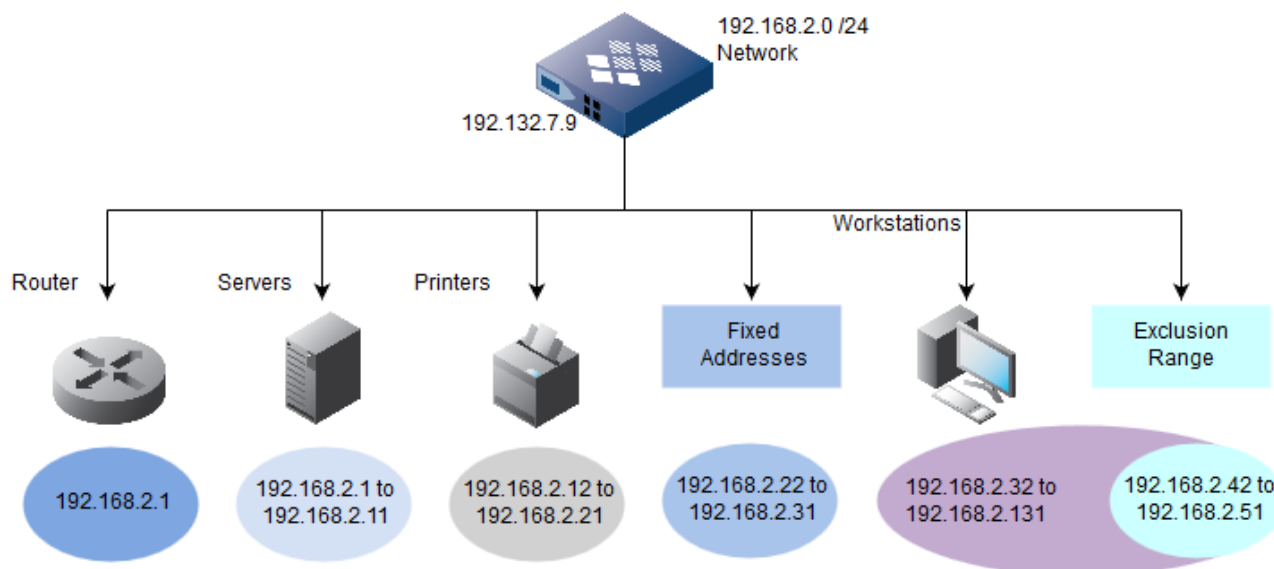
This example describes how to create a /24 network template and how to use the template to create a 192.168.2/24 network with the following configurations:

- First address 192.168.2.1 is reserved for the router
- Next 10 addresses (192.168.2.2 to 192.168.2.11) reserved for servers
- Next 10 addresses (192.168.2.12 to 192.168.2.21) reserved for printers
- Next 10 addresses (192.168.2.22 to 192.168.2.31) assigned as fixed addresses

- 100 addresses (192.168.2.32 to 192.168.2.131) reserved for workstations. The appliance assigns these dynamically.
- 10 addresses (192.168.2.42 to 192.168.2.51) are in an exclusion range. If you assigned static addresses to certain hosts in the middle of an address range template, you can exclude the addresses from the address range template so the appliance does not assign these IP addresses to clients

The below figure illustrates the configurations of the 192.168.2/24 network using the network template you create:

*Creating a Network Using a Template*



Use the following steps to create the sample network template (shown in the above figure).

1. Create the following DHCP range templates. For information, see [Adding IPv4 Range Templates](#) above.
  - Server template with the following values:
    - **Name:** Servers
    - **Offset:** 2
    - **Number of Addresses:** 10
    - **Comment:** Address range 2 to 11 for Servers
  - Printer template with the following values:
    - **Name:** Printers
    - **Offset:** 12
    - **Number of Addresses:** 10
    - **Comment:** Address range 12 to 21 for printers.
  - Workstation template with the following values:
    - **Name:** Workstations
    - **Offset:** 32
    - **Number of Addresses:** 100
    - **Comment:** Address range 32 to 131 for DHCP on workstations
  - Exclusion range with the following values. You must modify the *Workstations* template to add the exclusion range. For information, see [Modifying IPv4 Range Templates](#) above.
    - **Name:** Exclusion
    - **Offset:** 42
    - **Number of Addresses:** 10
    - **Comment:** Excluding addresses 42 to 51 from the DHCP range 32 to 131.
2. Create a fixed address/reservation template with the following values. For information, see [Adding IPv4 Fixed Address/Reservation Templates](#) above.
  - **Name:** Router
  - **Comment:** Fixed address template
  - **Offset:** 1
  - **Number of Addresses:** 1



3. Create a fixed address/reservation template with the following values. For information, see [Adding IPv4 Fixed Address/Reservation Templates](#) above.
  - **Name:** myFixedAddress
  - **Comment:** Fixed address template
  - **Offset:** 22
  - **Number of Addresses:** 10
4. Create a network template with the following values. For information, see [Adding IPv4 Network Templates](#) above.
  - **Name:** myNetworkTemplate
  - **Netmask:** Select /24 as the fixed subnet mask for the network
  - **Comment:** Network template for /24 network
  - **Automatically create an inverse-mapping zone:** Select this so that the NIOS appliance automatically creates the corresponding reverse-mapping zone for the network.
5. Add the DHCP range templates *Servers*, *Printers*, and *Workstations* to the network template.
6. Add the fixed address/reservation template *myFixedAddress* to the network template.
7. Add a fixed address with the following values:
8. Create a network using the network template *myNetworkTemplate* with the following values. For information, see [Adding IPv4 Networks](#).
  - **Address:** Enter the IP address 192.168.2.0 of the network that you want to create using the template.
  - **Select template:** Select the network template *myNetworkTemplate*.
9. To verify your configuration, from the **Data Management** tab, select the **DHCP** tab -> **Templates** tab. Select *myNetworkTemplate* and click the Edit icon. In the *Network Template* editor, click the **Templates** tab. The Grid Manager displays the DHCP range templates and fixed address templates.
10. Click **Restart** to restart services.

## About IPv6 DHCP Templates

You can use templates to create DHCP IPv6 ranges, fixed addresses, roaming hosts, and networks. You can create the following IPv6 templates:

- A DHCP range template that specifies an offset and the total number of addresses in a range. You can add a DHCP range template to a network template. For more information, see [About IPv6 Range Templates](#).
- A fixed address template, containing information for creating fixed addresses and roaming hosts. You can add a fixed address template to a network template. For information, see [About IPv6 Fixed Address Templates](#).
- A network template, containing basic network properties for creating networks. It is also a container that holds your DHCP range templates and fixed address/reservation templates. When you create a network using a network template, the network inherits the properties of the range and fixed address/reservation templates. You can create a network in any network view using a network template. For information, see [Adding IPv6 Network Templates](#).

Because you can potentially add DHCP range and fixed address/reservation templates to a network template, create the DHCP range and fixed address/reservation templates before you create a network template.

## About IPv6 Range Templates

You can create range templates to specify an offset and the number of addresses allocated to a range. Note that you cannot create templates for prefix-delegated ranges because the start or end prefix can be outside of the subnet address boundary.

After you create a DHCP range template, you can configure additional properties such as exclusion ranges and DHCP properties, as described in [Modifying IPv6 Range Templates](#) below. Then when you use the template to create a DHCP range, the range inherits the properties of the template. You can also include a DHCP range template in a network template to automatically create a DHCP range when you use that network template.

If you have deployed the Cloud Network Automation license on the Grid Master, you can configure range templates for cloud delegation. If you select a default Cloud Platform Appliance for a template, all ranges you create using this template will delegate authority to the same Cloud member. Note that when a Cloud member is removed from the Grid, the delegation will also be removed from the template. For information about Cloud Network Automation, see [Deploying Cloud Network Automation](#).

 **Note**

Infoblox does not support global IPv6 prefix delegation for IPv6 range templates.

## Adding IPv6 Range Templates

To create an IPv6 range template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab.
2. Click the drop-down menu of the Add icon and select **IPv6 Range Template**.
3. In the *Add IPv6 Range Template* wizard, complete the following:
  - **Name:** Enter a name that helps identify the IPv6 DHCP range template.
  - **Offset:** An offset in a DHCP range template determines the starting IP address of the range. The appliance adds the offset value you enter here to the start IP address of the network in which you create a DHCP range using this template. That IP address becomes the start IP address of the DHCP range. For example, you specify an offset value of 10 for the 2001:db8:1263:/48 network using the DHCP range template, the appliance creates a range with the start address 2001:db8:1263:0:0:0:a.
  - **Number of Addresses:** Enter the total number of IPv6 addresses to be included in the DHCP range.
  - **Comment:** Optionally, enter additional information about the template.

The **Cloud** section appears when the Cloud Network Automation license is installed on the Grid Master. For information, see [Deploying Cloud Network Automation](#). To configure this template for cloud delegation, complete the following:

**Use for cloud delegation:** Select this checkbox to enable cloud delegation for this template.

### Delegate authority from the Grid Master

- **Delegate To:** In a non-cloud API request, this parameter defines the default member to which authority is delegated. In a cloud API request, the appliance ignores this parameter, which allows you to use this template to create an object on different Cloud Platform Appliances. Click **Select** to choose the default Cloud Platform Appliance to which you want to delegate authority. The *Member Selector* displays only Cloud Platform Appliances in the Grid. Click the member, and Grid Manager displays the member name next to this field.
4. Click **Next** and select one of the following to provide DHCP services for the range:
    - **None (Reserved Range):** Select this if you want to reserve this address range for static hosts. Addresses in this range cannot be allocated as dynamic addresses. You can allocate the next available IP from this range to a static host. This is selected by default.
    - **Grid Member:** Click **Select** and choose a Grid member from the drop-down list.
  5. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attribute](#).
  6. Save the configuration.

## Modifying IPv6 Range Templates

You can modify the properties of a DHCP range template and define an exclusion range. For the exclusion range in the template, the start and end addresses are determined by the number of offsets in the DHCP range template's start address and the number of IP addresses in the exclusion range. For more information about exclusion ranges, see [About DHCP Ranges](#).

To modify a DHCP range:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* checkbox, and then click the Edit icon.
2. The *IPv6 DHCP Range Template* editor contains the following tabs from which you can modify data:
  - **General:** Modify general information as described in Adding IPv6 Range Templates above.
  - **Member Assignment:** Change the Grid member that provides DHCP services for ranges created from this template. For information, see Adding IPv6 Range Templates above.

- **Extensible Attributes:** Add and delete extensible attributes that are associated with this template. You can also modify the values of the extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [About Administrative Permissions](#).
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify data:
    - **Exclusion Ranges:** Configure a range of IP addresses that the appliance does not use for dynamic address assignments. **Exclusion Ranges:** Configure a range of IP addresses that the appliance does not use for dynamic address assignments. Complete the following:
      - **Offset:** An offset for an exclusion range determines the start IP address of the exclusion range. The appliance adds the offset value you enter here to the start IP address of the DHCP range created using this template. That IP address becomes the start IP address of the exclusion range.
      - **Number of Addresses:** Enter the number of IP addresses to be included in the exclusion range.
      - **Comment:** Enter useful information about the exclusion range. Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
  4. Save the configuration.

## About IPv6 Fixed Address Templates

A fixed address template is useful when you want to create multiple fixed addresses in a network. When you create a fixed address template, you specify the offset value and number of fixed addresses to be created. You can also specify additional properties for the fixed addresses.

Note that you can use the template to create address-based fixed addresses. You cannot specify prefixes in the template because a fixed address could use a prefix that is not part of the subnet to which the fixed address belongs. You can enter prefixes when you create the individual fixed address objects using the template.

If you have deployed the Cloud Network Automation license on the Grid Master, you can configure fixed address templates for cloud delegation. When you configure a template for cloud delegation, all fixed addresses you create using this template will inherit authority delegations from their parent objects. For information about Cloud Network Automation, see [Deploying Cloud Network Automation](#).

### Adding IPv6 Fixed Address Templates

To create an IPv6 fixed address template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab.
2. Click the drop-down menu of the Add icon and select **IPv6 Fixed Address Template**.
3. In the *Add IPv6 Fixed Address Template* wizard, enter the following:
  - **Name:** Enter a name that helps identify the IPv6 fixed address template. For example, you can enter **HP Printer** when you create a template that contains settings for assigning fixed addresses or reservations to HP printers.
  - **Comment:** Optionally, enter additional information about the template.
  - **Use for cloud delegation:** When you select this checkbox, all fixed addresses you create using this template inherit authority delegation from their parent objects. In the **Optional Settings For Range of Objects** section, do the following:
    - **Offset:** An offset in a fixed address template determines the IP address of the first fixed address created from the template. The appliance adds the offset value you enter here to the start IP address of the network in which you create objects using this template, and that IP address becomes the IP address of the object. For example, you specify an offset value of 50 for the 2001:db8:1263:/48 network, when you create a fixed address using the fixed address template, the appliance assigns it the address 2001:db8:1263:0:0:0:0:32.
    - **Number of Addresses:** Enter the number of IP addresses to be used as fixed addresses or roaming hosts. Note that the appliance uses the offset and number of addresses only when this template is used in a network template.
4. Click **Next** to configure or override DHCP options as described in [Defining General IPv6 Properties](#).
5. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).
6. Save the configuration.

## Modifying IPv6 Fixed Address Templates

To modify a fixed address template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* checkbox, and then click the Edit icon.
2. The *IPv6 Fixed Address Template* editor contains the following tabs from which you can modify data:
  - **General**: Modify general information for the template as described in Adding IPv6 Fixed Address Templates above.
  - **IPv6 DHCP Options**: Keep the inherited DHCP options or override them and enter unique settings for the template. For information, see [Defining General IPv6 Properties](#).
  - **Extensible Attributes**: Add and delete extensible attributes that are associated with the template. You can also modify the values of the extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions**: This tab appears only if you belong to a superuser admin group. For information, see [About Administrative Permissions](#).
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify data:
  - **IPv6 DDNS**: Keep the inherited DDNS settings or override them and enter unique settings for the template. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).
  - **Filters**: You can keep the inherited IPv6 logic filters or override them and add a new IPv6 logic filter. For information, see [Applying Filters to DHCP Objects](#).  
Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
4. Save the configuration.

## About IPv6 Network Templates

You can create IPv6 network templates to facilitate network configuration. You can use network templates to create networks in any network view. When you create a network template, you do not specify a network address. You enter the network address when you create an actual network from the template. You can specify a netmask or allow the user to define the netmask when they create the actual network.

A network template is useful for setting up a network with fixed addresses and DHCP ranges already defined. You can add DHCP range or fixed address templates to a network template.

### Note

You cannot configure the following DHCP options in an IPv6 network template: server-id (Option 2), preference (option 7), and unicast (Option 12). These options are valid only for a DHCP member.

When you enable support for RIR updates, you can create IPv6 network templates specific for RIR associated networks. For information about RIR updates, see [RIR Registration Updates](#).

You can also configure network templates for cloud delegation if you have deployed the Cloud Network Automation on the Grid Master. If you select a default Cloud Platform Appliance for a template, all networks you create using this template will delegate authority to the same Cloud member. Note that when a Cloud member is removed from the Grid, the delegation will also be removed from the template. For information about Cloud Network Automation, see [Deploying Cloud Network Automation](#).

## Adding IPv6 Network Templates

To create a network template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab.
2. Click the drop-down menu of the **Add** icon and select **IPv6 Network Template**.
3. In the *Add IPv6 Network Template* wizard, do the following:
  - **RegionallInternetRegistry**: This section appears only when support for RIR updates is enabled. For information about RIR, see [RIR Registration Updates](#). Complete the following to create a network template for an RIR IPv6 network container or network:

- **Internet Registry:** Select the RIR from the drop-down list. The default is **RIPE**. When you select **None**, the network is not associated with an RIR organization.
- **Organization ID:** Click **Select Organization** and select an organization from the *RIR Organization Selector* dialog box.
- **Registration Status:** The default is **Not Registered**. When using this template to add an RIR allocated network, you can change this to **Registered** and select the **Do not update registrations** checkbox below. Note that when you select **API** as the communication method, the registration status will be updated automatically after the registration update is completed. However, when you select **Email** as the communication method, the registration status will not be automatically updated. If you are creating a new network and the registration update is completed successfully, the status will be changed to **Registered**. If the update fails, the status will be changed to **Not Registered**.
- **Registration Action:** Select the registration action from the drop-down list. When you select **Create**, the appliance creates the IPv4 network and assigns it to the selected organization. When you select **None**, the appliance does not send registration updates to RIPE. When you use this template to add an existing RIR allocated network to NIOS, select **None**. When you use this template to add networks to an RIR allocated network (a parent network), select **Create**. Ensure that the parent network associated with an RIR organization already exists.
- **Do not update registrations:** Select this checkbox if you do not want the appliance to submit RIR updates to RIPE. By default, the appliance sends updates to the RIR database based on the configured communication method.
- **IPv6 Prefix:** If you are adding a template for a previously defined global IPv6 prefix, you can select it from the drop-down list.
- **Name:** Enter a name that helps identify the network template.
- **Netmask:** Select one of the following options:
  - **Fixed:** Select this and adjust the netmask slider to a fixed netmask for this network template. When you select this option, users cannot specify another netmask when they use this template to create a network. For example, if you select /24 as the fixed netmask, all networks created using this template have a /24 netmask. The slider moves to the CIDR value associated with the selected prefix when you choose a global IPv6 prefix.
  - **Allow User to Specify Netmask:** Select this to allow users to specify the subnet mask when creating networks using this template.
- **Comment:** Enter useful information about the template.
- **Automatically create a reverse-mapping zone:** This function is enabled if the fixed netmask of the template is a multiple of 4 (4, 8, 24, and so on), or if you select the **Allow User to Specify Netmask** option. Select this if you want the appliance to automatically create the corresponding reverse-mapping zone for the networks created using this template. These zones are created in the DNS view assigned to receive dynamic DNS updates at the network level.

The **Cloud** section appears when the Cloud Network Automation license is installed on the Grid Master. For information, see [Deploying Cloud Network Automation](#). To configure this template for cloud delegation, complete the following:

**Use for cloud delegation:** Select this checkbox to enable cloud delegation for this template.

#### Delegate authority from the Grid Master

**Delegate To:** In a non-cloud API request, this parameter defines the default member to which authority is delegated. In a cloud API request, the appliance ignores this parameter, which allows you to use this template to create an object on different Cloud Platform Appliances. Click **Select** to choose the default Cloud Platform Appliance to which you want to delegate authority. The *Member Selector* displays only Cloud Platform Appliances in the Grid. Click the member, and Grid Manager displays the member name next to this field.

4. Click **Next** to assign Grid members to this network template. Ensure that you include members that are associated with other templates that you plan to add to this network template. You can assign one or multiple members to this template. However, you cannot assign a combination of NIOS Grid members and vNIOS Grid members to the template.
  - Click the Add icon to add a Grid member as a DHCP server for the networks created using this template. Select the Grid member from the *Member Selector* dialog box. Keep in mind, DHCP properties for the

network are inherited from this member. Networks created using this template can be served by multiple members, but a member can serve networks in one network view only.

5. Click **Next** to associate Active Directory Sites with the network. For more information, see [Associating Active Directory Sites with Networks](#).
6. Click **Next**, and then click the Add icon to include DHCP range and fixed address templates in the network template. Choose the template that you want to include in this network template. Use SHIFT+click and CTRL+click to select multiple templates.  
You can remove a template from the list by selecting the template and clicking the Delete icon.
7. Click **Next** to configure or override DHCP options as described in [Defining General IPv6 Properties](#).
8. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).  
If you are adding an RIR network, the RIR network attribute table appears. For information about these attributes and how to enter them, see [RIR Network Attributes](#).
9. Save the configuration.

## Modifying IPv6 Network Templates

To modify and set the properties of a network template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* checkbox, and then click the Edit icon.
2. The *Network Template* editor contains the following tabs from which you can modify data:
  - **General**: Modify general information described in Adding IPv6 Network Templates above.
  - **Member Assignment**: Change the Grid members that provide DHCP services for networks created from this template. For information, see [Adding IPv6 Networks](#).
  - **Templates**: Add or delete DHCP range and fixed address templates. For information, see [Adding IPv6 Range Templates](#) and [Adding IPv6 Fixed Address Templates](#).
  - **IPv6 DHCP Options**: Keep the inherited DHCP options or override them and enter unique settings for the template. For information, see [Defining General IPv6 Properties](#).
  - **RIR Registration**: Modify RIR network information. This tab appears only when support for RIR updates is enabled. For information, see [Modifying RIR Network Data](#).
  - **Extensible Attributes**: Add and delete extensible attributes that are associated with the template. You can also modify the values of the extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions**: This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#).
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify data:
  - **IPv6 DDNS**: Keep the inherited DDNS settings or override them and enter unique settings for the template. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).
  - **Filters**: You can keep the inherited IPv6 logic filters or override them and add a new IPv6 logic filter. For information, see [Applying Filters to DHCP Objects](#).  
Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
4. Save the configuration.

## Viewing and Deleting Templates

You can view a list of all IPv4 and IPv6 DHCP templates and delete the templates that are not required.

### Viewing Templates

To view a list of all IPv4 and IPv6 DHCP templates:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab.
2. Grid Manager displays the following information:
  - **Name**: The name of the template.
  - **Type**: The template type, such as **IPv4 Network Template** or **IPv6 Network Template**.
  - **Comment**: The information you entered about the template.



- **Site:** The site to which the template belongs. This is one of the predefined extensible attributes.

You can select predefined and user defined extensible attributes for display. You can also do the following in this panel:

- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).
- Sort the displayed data in ascending or descending order by column.
- Delete a selected template or multiple templates. For information, see [Deleting Templates](#) below.
- Use filters and the **Goto** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Goto** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Select an object and edit its information.
- Print or export the data in the panel.

## Deleting Templates

To delete a template:

1. From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> *template* checkbox, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.

## Managing IPv6 DHCP Data

This section explains how to configure and manage IPv6 DHCP data. It contains the following topics:

- [Configuring IPv6 Networks](#)
- [Defining Global IPv6 Prefixes](#)
- [Managing IPv6 Networks](#)
- [Configuring IPv6 Shared Networks](#)
- [Configuring IPv6 Fixed Addresses](#)
- [Configuring IPv6 Address Ranges](#)
- [Viewing IPv6 DHCP Objects](#)
- [Setting the Prefix Length Mode for DHCPv6](#)

## Configuring IPv6 Networks

To configure DHCP services for an IPv6 network and the resources in the network, perform the following tasks:

1. To facilitate network creation, you can specify the IPv6 global prefixes for the Grid. For more information, see [Defining Global IPv6 Prefixes](#).
2. Create a network and assign it to Grid members. For information, see [Managing IPv6 Networks](#) and [About IPv6 Shared Networks](#).
3. Optionally, configure DHCP properties for the network. You can override properties set at the Grid or member level and enter unique values for the network and fixed addresses. For information, see [Configuring DHCPv6 Properties](#) and [Configuring DHCP IPv4 and IPv6 Common Properties](#).
4. Optionally, assign zones to a network. For information, see [Associating Networks with Zones](#).
5. Add a DHCP range to the network and assign it to a member. For information, see [Configuring IPv6 Address Ranges](#).
6. Optionally, add exclusions to the DHCP range for addresses that are not used for dynamic allocation. For information, see [Modifying IPv6 Address Ranges](#).
7. Optionally, configure DHCP properties for the address range. You can override properties set at an upper level and enter unique values for the address range. For information, see [Modifying IPv6 Address Ranges](#).



8. Optionally, add fixed addresses to the network and configure DHCP properties for them. A fixed address may also be associated with a device port through a Port Reservation. For information, see [Configuring IPv6 Fixed Addresses](#).
9. Start the DHCP service and the IPv6 DHCP service. For more information, see [Starting DHCP Services on a Member](#).

## Defining Global IPv6 Prefixes

To simplify network creation, you can define IPv6 prefixes that are used for networks served by the Grid members. If your organization is assigned IPv6 prefixes, you can enter them globally at the Grid level, and then just select the appropriate IPv6 prefix when you define the network and network templates. You can create multiple global prefixes. When you define an IPv6 network and network templates, you must adjust the slider to the desired netmask as per the CIDR in the prefix. Use the netmask slider to select /64 as the CIDR, if an IPv6 prefix is unavailable for an IPv6 network.

To add global IPv6 prefixes:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Expand the Toolbar and click **Grid DHCP Properties**.
3. In the *Grid DHCP Properties* editor, select the **IPv6 Global Prefixes** tab.
4. Click the Add icon and enter a name for the prefix and the prefix. Select the **Default** checkbox to specify a default IPv6 prefix for the Grid.
5. Save the configuration.

## Managing IPv6 Networks

You can create an IPv6 network from scratch or create a network template and then use that template to create one or more networks. Using a network template facilitates creating multiple IPv6 networks with similar properties. You can also create an IPv6 network from the Tasks Dashboard. For information about the Tasks Dashboard, see [Tasks Dashboard](#). An IPv6 network inherits its DHCP options & DDNS settings from its shared network, if it is part of a shared network, or from the member to which it is assigned.

If you have enabled support for RIR (Regional Internet Registry) updates and are adding an RIR IPv6 network container or network to NIOS, Grid Manager displays an RIR section in the *Add IPv6 Network* wizard. You must enter RIR related information in this section in order for NIOS to associate the newly added network with an RIR organization. For more information about RIR address allocation and updates, see [RIR Registration Updates](#).

## Adding IPv6 Networks

To add an IPv6 network:

1. Select the **Data Management** tab.
2. If you have more than one network view in the system, select the network view in which you want to add the network.
3. Select the **DHCP** tab -> **Networks** tab.
4. In the **Networks** section, expand the **Toolbar** and click the **Add** drop-down list and select **Network-> IPv6** or click the Add icon drop-down list and select **IPv6 Network**.
5. In the *Add IPv6 Network* wizard, select one of the following and click **Next**:
  - **Add IPv6 Network**: Click this to add an IPv6 network from scratch.
  - **Add IPv6 Network using Template**: To use a template, click this, and then click **Select Template** and select an IPv6 network template. For information about network templates, see [About IPv6 Network Templates](#). When you use a template to create a network, the configurations of the template apply to the new network. The appliance populates the template properties in the wizard when you click **Next**. You can then edit the pre-populated properties. If the template specified a fixed netmask, you cannot edit the netmask.
6. Complete the following and click **Next**:
  - **Regional Internet Registry**: This section appears only when support for RIR updates is enabled. For information about RIR, see [RIR Registration Updates](#). Complete the following to create an RIR IPv6 network container or network:

- **Internet Registry:** Select the RIR from the drop-down list. The default is **RIPE**. When you select **None**, the network is not associated with an RIR organization.
- **Organization ID:** Click **Select Organization** and select an organization from the *RIR Organization Selector* dialog box.
- **Registration Status:** The default is **Not Registered**. When adding an RIR allocated network, you can change this to **Registered** and select the **Do not update registrations** checkbox below. Note that when you select **API** as the communication method, the registration status will be updated automatically after the registration update is completed. However, when you select **Email** as the communication method, the registration status will not be automatically updated. If you are creating a new network and the registration update is completed successfully, the status will be changed to **Registered**. If the update fails, the status will be changed to **Not Registered**. The updated status and timestamp are displayed in the **Status of last update** field in the *IPv6 Network Container* or *IPv6 Network* editor.
- **Registration Action:** Select the registration action from the drop-down list. When you select **Create**, the appliance creates the IPv4 network and assigns it to the selected organization. When you select **None**, the appliance does not send registration updates to RIPE. When you are adding an existing RIR allocated network to NIOS, select **None**. When you are adding networks to an RIR allocated network (a parent network), select **Create**. Ensure that the parent network associated with an RIR organization already exists.
- **Do not update registrations:** Select this checkbox if you do not want the appliance to submit RIR updates to RIPE. By default, the appliance sends updates to the RIR database based on the configured communication method.
- **Netmask:** Use the netmask slider to select the appropriate number of subnet mask bits for the network. Select /64 as the CIDR, if an IPv6 prefix is unavailable for an IPv6 network. When the prefix is available, you must adjust the slider to the desired netmask as per the CIDR in the prefix.
- **Networks:** Do one of the following to add new networks:
  - Click the Add icon to enter a new network. If you are adding a network for a previously defined global IPv6 prefix, you can select the prefix from the **IPv6 Prefix** drop-down list. The default is **None**, which means that you are not creating an IPv6 network for a previously defined subnet route. If you have defined a global prefix at the Grid level, the default is the global prefix value. Click **Add** and Grid Manager adds a row to the table. Enter the network address in the **Network** field. When you enter an IPv6 address, you can use double colons to compress a contiguous sequence of zeros. You can also omit any leading zeros in a four-hexadecimal group. For example, the complete IPv6 address 2001:0db8:0000:0000:0000:0000:0102:0304 can be shortened to 2001:db8::0102:0304. Note that if there are multiple noncontiguous groups of zeros, the double colon can only be used for one group to avoid ambiguity. The appliance displays an IPv6 address in its shortened form, regardless of its form when it was entered. Click **Add** again to add another network. You can also select a network and click the **Delete** icon to delete it.
  - or
  - Click the Next Available icon to have the appliance search for the next available network. Complete the following in the Next Available Networks section:
    - **Create new network(s) under:** Enter the network container in which you want to create the new network. When you enter a network that does not exist, the appliance adds it as a network container. When you enter a network that is part of a parent network, the parent network is converted into a network container if it does not have a member assignment or does not contain fixed addresses and host records that are served by DHCP. You can also click **Select Network** to select a specific network in the *Network Selector* dialog box. For information about how the appliance searches for the next available network, see [Obtaining the Next Available](#).
    - **Number of new networks:** Enter the number of networks you want to add to the selected network container. Note that if there is not enough network space in the selected network to create the number of networks specified here, Grid Manager displays an error message. The maximum number is 20 at a time. Note that when you have existing networks in the table and you select one, the number you enter here includes the selected network.
    - Click **Add Next** to add the networks. Grid Manager lists the networks in the table. You can click **Cancel** to reset the values.

You must click **Add Next** to add the network container you enter in the Next Available

Networks section. If you enter a network in the Next Available Networks section and then use the Add icon to add another network, the appliance does not save the network you enter in the Next Available Networks section until you click **Add Next**.

- **Comment:** Enter additional information about the network, such as the name of the organization it serves.
  - **Automatically Create Reverse-Mapping Zone:** This function is enabled if the netmask of the network is a multiple of four, such as 4, 8, 12 or 16. Select this to have the appliance automatically create reverse-mapping zones for the network. A reverse-mapping zone is an area of network space for which one or more name servers have the responsibility for responding to address-to-name queries. These zones are created in the DNS view assigned to receive dynamic DNS updates at the network view level.
  - **Disable for DHCP:** Select this if you do not want the DHCP server to provide DHCP services for this network at this time. This feature is useful when you are in the process of setting up the DHCP server. Clear this after you have configured the server and are ready to have it serve DHCP for this network. Note that disabling an IPv6 network may take a longer time to complete depending on the size of the data.
  - The **Cloud** section appears when the Cloud Network Automation license is installed on the Grid Master. For information, see [Deploying Cloud Network Automation](#). To delegate authority for this network, complete the following:
    - **Delegate authority from the Grid Master**
    - **Delegate To:** This field indicates whether the authority for the network you want to create has already been delegated to a Cloud Platform Appliance. Click **Select** to choose the Cloud Platform Appliance to which you want to delegate authority. The *Member Selector* displays only Cloud Platform Appliances in the Grid. Click the member, and Grid Manager displays the member name next to this field. This cloud member now assumes authority for this network, and the Grid Master does not have authority any more. You can also click **Clear** to remove authority delegation from the selected Cloud Platform Appliance and return authority back to the Grid Master.
7. Click **Next** and add one or more Grid members as DHCP servers for the network.
    - Click the **Add** icon and select a Grid member from the *Member Selector* dialog box. Note that, some DHCP properties for the network are inherited from this member. The network can be served by multiple members, but a member can serve networks in one network view only.
  8. Click **Next** to associate Active Directory Sites with the network. For more information, see [Associating Active Directory Sites with Networks](#).
  9. (*Applies only to Network Insight*) Click **Next** to initiate or disable discovery of the new network(s). Discovery settings differ based on whether you are defining one network or multiple networks.
    - **Configuring one network:** discovery settings include the following: **Enable Discovery** and **Immediate Discovery**, selecting a Probe member to perform the discovery; and **Polling Options**, which define how the network will be discovered by the Probe member. By default, all Polling Options discovery settings are inherited from the parent network (or Grid, if no parent exists) unless you click **Override**. Polling Options govern the protocols used to query and collect information about the network devices being discovered.  
or
    - **Configuring more than one network:** If the networks are child networks, they automatically inherit the settings of the parent network, including discovery settings and the discovery member. These settings will not appear in the wizard page. For discovery of multiple networks, you can only enable or disable **Immediate Discovery**. Click **Next** to override the DHCP properties described in [Defining General IPv6 Properties](#).
  10. Assign VLAN objects to the network. For more information, see [VLAN Management](#).
  11. As part of creating a network, you can provision the network on an actual device (switch, router, or switch-router), that is discovered and managed through the Grid Manager.
    - Begin by checking the **Enable Network Provisioning** checkbox, and clicking the **Select Device** button. Choose your device from the Device Selector dialog. (Click **Clear** to remove the setting. For more information, see the section [Using the Device Selector](#).)
    - If you performed DHCP configuration in the previous step of the Add Network Wizard, the **Router IP** value will automatically be populated with the DHCP Router IP address value. Otherwise, you enter the standard router IP address.
    - If required for the newly provisioned network to ensure that attached devices receive DHCP auto-configuration, enable the **DHCP Forwarding** checkbox. For this setting, if a DHCP Failover was previously configured, the IP addresses defined for DHCP failover are automatically used for the DHCP forwarding configuration.
    - You will also need to choose an interface on the selected device on which to provision the network by selecting it from the **Interface** drop-down menu. Grid Manager ensures that only those interfaces that can

support provisioning, and are available for provisioning (that do not have an **Operation Status** of Up), appear in the drop-down menu.

- Otherwise, when creating networks and provisioning them on managed devices, you can create a VLAN on which to provision the network by clicking the **Create VLAN** option and entering the **VLAN Name** and **VLAN ID**. Ensure that the **VLAN ID** value you enter is appropriate for the application - don't create a new VLAN and provision a network for a VLAN value that is already actively carrying traffic for another routing domain.

If a selected device does not support VLANs, the **Create VLAN** option will not appear.

12. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Managing Extensible Attributes](#).

If you are adding an RIR network, the RIR network attribute table appears. For information about these attributes and how to enter them, see [RIR Network Attributes](#). You can preview the information before the appliance submits updates to the RIPE database. To preview registration updates, click **Preview RIR Submissions**. For more information, see [Previewing Registration Updates](#).

Note that you cannot leave an optional RIR attribute value empty. If you do not have a value for an RIR attribute, you must delete it from the table. You can enter up to 256 characters for all RIR attributes.

You need to assign a **Subscriber Member Site** to add subscriber service related extensible attributes in order to populate Subscriber Cache.

13. As the final step in the Add IPv6 Network wizard, you define when Grid Manager creates the new network by scheduling it. You also schedule when the associated port control task executes (if a port configuration has been specified).
  - To create the new network and its associated port configuration immediately, select **Now**. Grid Manager synchronizes the port control task to take place at the same time as the creation of the new network.
  - You can choose to have Grid Manager execute the port control task at a later time by selecting **Later**. Choose a **Selected time** by entering or selecting a **Start Date** (click the calendar icon to choose a calendar date) and a **Start Time**, and choose a **Time Zone**.
14. Choose one of the following from the **Save &...** drop-down button menu:
  - Click **Save & Close** to add the new network and close the wizard (this is the default).
  - Click **Save & Edit** to add the new network and launch the editor.
  - Click **Save & New** to add the new network and launch the wizard again to add another network.

At any step during the wizard, you can click Schedule for Later to schedule the task. In the Schedule Change panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

After you create a network, you can do the following:

- Add it to a shared network. For more information, see [Adding IPv6 Shared Networks](#).
- Use the split network feature to create subnets for the network. For information, see [Splitting IPv6 Networks into Subnets](#).
- Use the join networks feature to create a parent network that encompasses multiple subnets into a larger network. For information, see [Joining IPv6 Networks](#). You can also create a shared network for subnets that are on the same network segment.
- View a list of networks. For more information, see [Viewing Networks](#).

## Modifying IPv6 Networks

You can modify existing network settings and override the Grid or member DHCP properties, with the exception of the network address and netmask.

To modify an IPv6 network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* checkbox, and then click the Edit icon.
2. The *IPv6 Network* editor contains the following basic tabs from which you can modify data:
  - **General Basic**: You can modify the following fields:
    - **Comment**: The information you entered for the network.

- **Disabled:** This field is displayed only if the selected network is a network without a child network under it. You can disable and enable existing networks instead of removing them from the database, if the selected network does not have a child subnet. This feature is especially helpful when you have to move or repair the server for a particular network. Note that disabling an IPv6 network may take a longer time to complete depending on the size of the data. When the Cloud Network Automation license is installed on the Grid Master, Grid Manager displays the following information in the **Cloud** section: **Cloud Usage**, **Owned By**, and **Delegated To**.

You cannot modify these fields. For more information, see [Adding IPv6 Networks](#) above.

- **Member Assignment:** Add or delete a Grid member that provides DHCP services for this network.
  - **IPv6 DHCP Options:** Keep the inherited DHCP properties or override them and enter unique settings for the network. For information, see [Defining General IPv6 Properties](#).
  - **Discovery:** Checking the **Enable Discovery** checkbox informs NIOS to begin discovering the network after you click **Save and Close**. You manage discovery polling settings local to each network from this page. For a complete overview of features on this page, see [Discovering Devices and Networks](#) and its subsections.
  - **Discovery Exclusions:** IP Addresses and IP ranges can be locally excluded from discovery by clicking the Add icon and selecting **Add IP Address** or **Add IP Range**. These IP addresses or IP ranges are selected from within the chosen network. For related information, see [Excluding IP Addresses from Discovery](#) and its subsections.
  - **Discovery Blackout:** Define extended time periods and regularly scheduled times when discovery and/or port configuration tasks will not take place on a network. Editing a network under DHCP, blackout settings apply only to the specified network. You also specify the scheduled time when the blackout period begins, and the duration of the blackout period. By default, the network inherits its discovery blackout settings from the Grid level. For related information, see [Defining Blackout Periods](#) and its subsections. Discovery blackout settings also can be defined for DHCP ranges.
  - **RIR Registration:** Modify RIR network information. This tab appears only when support for RIR updates is enabled. For information, see [Modifying RIR Network Data](#).
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of the extensible attributes. For information, see [Managing Extensible Attributes](#).
  - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#).
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
    - **General Advanced:** You can associate zones with a network. For information, see [Associating Networks with Zones](#).
    - **IPv6 DDNS:** Keep the inherited DDNS settings or override them and enter unique settings for the network. Note that you must click **Override** and select **Enable DDNS updates** for the DDNS settings you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).
    - **Filters:** You can keep the inherited IPv6 logic filters or override them and add a new IPv6 logic filter. For information, see [Applying Filters to DHCP Objects](#).

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.
  4. Save the configuration or click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Deleting IPv6 Networks

When you delete a network, all of its data, including all DHCP records, subnets, and records in its subnets, is deleted from the database. Because of the potentially large loss of data that can occur when you delete a network, the appliance stores the deleted network in the Recycle Bin. You can restore a deleted network from the Recycle Bin, if enabled. You can also disable a network instead of deleting it. For information, see [Modifying IPv6 Networks](#) above.

To delete a network:



1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> *network* checkbox, and then select **Delete** or **Schedule Delete** from the Delete drop-down menu.
2. To delete the network now, in the Delete Confirmation dialog box, click **Yes**. Grid Manager displays a warning message. Click **Yes** to continue or **No** to cancel the process. To schedule the deletion, see [Scheduling Deletions](#). You can also choose to export all network data to a CSV file before deleting. To do this, click the **Export & Delete** button. For more information about this option, see [Configuring IPv4 Networks](#).

The appliance puts the deleted network in the Recycle Bin, if enabled. Click **Restore** in the Recycle Bin to recover the deleted data. Click **Yes** in the Restore Item dialog box to restore or **No** to cancel the process. Note that deleting and restoring an IPv6 network may take a longer time to complete depending on the size of the data.



#### Note

You cannot delete a network that has a VLAN object assigned to it. For more information, see [Assigning VLANs to a Network](#).

## Configuring IPv6 Shared Networks

You can combine two or more contiguous IPv6 networks into a shared network. When you do, the DHCP server allocates IP addresses from both subnets. To create a shared network, create the individual subnets, and then create the shared network and add the subnets to it. For more information about shared networks, see [About Shared Networks](#).

### Adding IPv6 Shared Networks

To add an IPv6 shared network:

1. Select the **Data Management** tab.
2. If you have more than one network view in the system, select the network view in which you want to add the network.
3. Select the **DHCP** tab -> **Networks** tab.
4. In the **Shared Networks** section, click the **Add** drop-down list and select **Shared Network-> IPv6** or or, click the Add icon drop-down list and select **IPv6 Shared Network**.
5. In the *Add IPv6 Shared Network* wizard, do the following:
  - **Name:** Enter the name of the shared network.
  - **Comment:** Enter information about the shared network.
  - **Disabled:** Select this if you want to enable the shared network at a later time. You can disable and enable existing networks instead of removing them from the database. This feature is especially helpful when you have to move or repair the server for a particular network.
6. Click **Next** and do the following to add networks:
  - a. Click the Add icon.
  - b. In the *Network Selector*, select the networks that you want to include in the shared network. Ensure that the networks are served by the same Grid members to avoid DHCP inconsistencies.
7. Click **Next** to configure DHCP properties described in [Defining General IPv6 Properties](#).
8. Click **Next** to enter values for required extensible attributes or add optional extensible attributes for the shared network. For information, see [Using Extensible Attributes](#).
9. Save the configuration or click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

For information on viewing shared networks, see [Viewing Shared Networks](#).

### Modifying IPv6 Shared Networks

To modify a shared network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Shared Networks** section -> *shared\_network* checkbox, and then click the Edit icon.

2. The *IPv6 Shared Network* editor contains the following tabs from which you can modify data:
  - **General**: Modify the fields **Name**, **Comments**, and **Disabled** as described in the previous section, [Adding IPv6 Shared Networks](#).
  - **Networks**: Displays the networks that are currently assigned to the shared network. You can add or delete a network. To add a network, click the Add icon. To delete a network, select the *network* checkbox, and then click the Delete icon.
  - **IPv6 DHCP Options**: Keep the inherited DHCP properties or override them and enter unique settings for the shared network. For information, see [Defining General IPv6 Properties](#).
  - **Extensible Attributes**: Add and delete extensible attributes that are associated with a specific network. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions**: This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#).
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
  - **IPv6 DDNS**: Keep the inherited DDNS settings or override them and enter unique settings for the shared network. Note that you must click **Override** and select **Enable DDNS updates** for the DDNS settings you configure in this tab to take effect. For information, see [Enabling DDNS for IPv4 and IPv6 DHCP Clients](#).
  - **Filters**: You can keep the inherited IPv6 logic filters or override them and add a new IPv6 logic filter. For information, see [Applying Filters to DHCP Objects](#).

Note that Grid Manager displays both the basic and advanced tabs the next time you log in to the GUI.

4. Save the configuration and click **Restart** if it appears at the top of the screen.  
or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Deleting IPv6 Shared Networks

Though you can delete the networks in a shared network, a shared network must have at least one network in it. To delete a shared network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Shared Networks** section -> *shared\_network* checkbox, and then select **Delete** or **Schedule Delete** from the drop-down menu.
2. To delete the shared network now, in the *Delete Confirmation* dialog box, click **Yes**. To schedule the deletion, see [Scheduling Deletions](#).

The appliance puts the deleted shared network in the Recycle Bin, if enabled.

## Configuring IPv6 Fixed Addresses

You can configure IPv6 fixed addresses with either an IPv6 address or prefix. You can assign prefix-based fixed addresses to routers so they can advertise the prefixes associated with a link. The fixed addresses also can be bound to interfaces on a network device, such as a switch or a router, that is discovered and managed under IPAM. DHCP hosts, in turn, use these prefixes to generate IP addresses using the stateless autoconfiguration mechanism defined in *RFC 2462, IPv6 Stateless Autoconfiguration*. You can also create IPv6 fixed addresses from the Tasks Dashboard. For information about the Tasks Dashboard, see [Tasks Dashboard](#).



### Note

IPv6 fixed addresses do not support dynamic DNS updates.



## Adding IPv6 Fixed Addresses



### Note

At any time during the wizard, you can click **Schedule for Later** to schedule the task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

To add an IPv6 fixed address:

1. Navigate to the network to which you want to add a fixed address, and then select **Fixed Address** from the Add drop down menu.  
or  
From any panel in the DHCP tab, expand the Toolbar and click **Add -> Fixed Address -> IPv6**.
2. In the *Add Fixed Address* wizard, select one of the following and click **Next**:
  - **Add IPv6 Fixed Address**  
or
  - **Add IPv6 Fixed Address Using Template**

Click **Select Template** and select the template that you want to use. When you use a template to create a fixed address, the configurations of the template apply to the new address. The appliance automatically populates the fixed address properties in the wizard. You can edit the pre-populated properties.

3. In this panel, the displayed network address can either be the last selected network or the network from which you are adding the fixed address. If no network address is displayed or if you want to specify a different network, click **Select Network**. When there are multiple networks, Grid Manager displays the *Select Network* dialog box. Specify one of the following:
  - Select **Address** to assign an IPv6 address to a fixed address. You can either enter an IPv6 address or select **Next Available IP** to obtain the next available IP address. Note that for Cloud Network Automation, Next Available IP is not available if the fixed address you want to create is within a delegated range.
  - Select **Prefix Delegation** to assign an IPv6 prefix. Enter the prefix and prefix length.
  - Select **Both** to assign an IPv6 prefix and address. Enter the IPv6 address, prefix, and prefix length.

Complete the following:

- **DUID**: Specify the DHCP Unique Identifier (DUID) of the DHCP client assigned to this fixed address.
- **Name**: Enter a name for the fixed address.
- **Comment**: Optionally, enter additional information.
- **Disabled**: Select this if you do not want the DHCP server to allocate this IP address at this time.

The **Cloud** section appears when the Cloud Network Automation license is installed on the Grid Master. For information, see [Deploying Cloud Network Automation](#). This section displays the following information:

- **Cloud Usage**: This field indicates whether this object is associated with any specific cloud extensible attributes or within a scope of delegation. It can be one of the following:
  - **Cloud from adapter**: Indicates that this object has been created by a cloud adapter and it may or may not be within a scope of delegation at the moment.
  - **Cloud from delegation**: Indicates that this object is within the scope of delegation or the object itself defines a scope of authority delegation, and it is not created by a cloud adapter.
  - **Used by cloud**: Indicates that this network or network container is associated with the extensible attribute **Is External** or **Is Shared** and the value is set to True, which implies the network is a private or shared network managed by the CMP, and it is not **Cloud from adapter** or **Cloud from delegation**.
  - **Non-cloud**: The object is a regular NIOS object and is not within the scope of any authority delegation nor is it associated with any of these extensible attributes: **Cloud API Owned**, **Is External** or **Is Shared**. NIOS admin users can modify this object based on their permissions.
- **Owned By**: A cloud object can be owned by the Grid Master or the cloud adapter. When the object is created by the Grid Master, this shows **Grid**. If the object is created by the cloud adapter, this shows **Adapter**.

#### Delegate authority from the Grid Master

- **Delegate To:** This field indicates whether the authority for the object you want to create has already been delegated. If so, it displays the name of the delegation.
4. Click **Next** to configure or override DHCP options as described in [Defining General IPv6 Properties](#).
  5. (*Applies only with Network Insight*) *This step is not required for creating a new Fixed Address.* In the current Wizard step, you can optionally define the following identification values and settings for the new object's port reservation:
    - Choose the **DeviceType:** **Router**, **Switch-Router**, **Switch**, **MSFT** (Microsoft) **Server**, **NetMRI**, **NIOS**, **VNIOS**, or **ESX** (VMware) **Server**.

The values on this page are not required for defining the actual port reservation in a later wizard step. Certain device types could be descriptively relevant based on the type of object you are creating. As an example, the **MSFT Server** designator helps identify the new object as a Microsoft Hyper-V Host. The **ESX Server** designator can be used to identify the new object as a VMware ESX Host. These values are not required and will not affect the functionality of the object.
    - Choose the **Device Vendor:** **Cisco**, **Juniper**, **Aruba**, **Dell**, **Infoblox**, or **HP**.
    - You can also enter a **Location** and a **Description**. These values are advisory and not required for configuration.

After you define this group of settings, you will still need to define a device port reservation, which is done in a later step.
  6. Click **Next** to initiate or disable discovery of the new Fixed Address. (*Applies only to Network Insight*) *This step is not required for creating a new Fixed Address.*
    - a. Choose either **Exclude from Network Discovery** or **Enable Immediate Discovery**. If you choose to Exclude, discovery will not execute on the Fixed Address. If you choose **Enable Immediate Discovery**, discovery will execute on the host after you save your settings. You may also choose to leave both options disabled.
    - b. By default, the new fixed address object inherits its SNMP credentials from those defined at the grid level. Should you wish to override them for a local set of credentials, check the **Override Credentials** checkbox and select the **SNMPv1/SNMPv2** or **SNMPv3** option and enter the locally used credentials.
    - c. You may also test the entered SNMP credentials by clicking **Test SNMP Credential**.

For descriptions of SNMP credentials for discovery, see the section [Configuring SNMP1/v2 Credentials for Polling](#) and [Configuring SNMPv3 Properties](#). These Grid-based values are inherited, by default, by each new object you create.

- For the new object, you can check the **Override CLI Credentials** checkbox to override the inherited set of CLI credentials taken from the Grid level. This set of credentials may be used for the device that is directly associated with the new object in its Port Reservation.
- You can also click **Test CLI Credentials** to enter and test a set of CLI login credentials against a device based on its IP address.

Port control operations require CLI credentials for the involved devices. (If you are not using port control for the new object, usage of CLI credentials is optional.) Because some IPAM and DHCP objects will use port control features as part of object creation, CLI credentials are automatically leveraged as part of discovery. Ensure you have the correct sets of CLI credentials for devices in your network. See the section [Configuring CLI Discovery Properties](#) for related information.

- SSH is the default for CLI operations. Check the **Allow Telnet** checkbox if you know the device involved in the object assignment may support Telnet but may not support SSH, or if you want Telnet as an option.

All port control operations require CLI credentials to be entered into Grid Manager. Because some IPAM and DHCP objects will use port control features as part of object creation, CLI credentials are automatically leveraged as part of discovery. Ensure you have the correct sets of CLI credentials for devices in your network.

7. Click **Next** to define the port reservation for the device port that will be associated with the new Fixed Address. *This step is not required for creating a new Fixed Address.* This feature set is also termed *port control* in the NIOS/Grid Manager system. The device to which the new Fixed Address will be associated should already be discovered and managed from the Grid Manager.
  - Begin by checking the **Reserve Port** checkbox. Note that reserving a switch port does not guarantee its availability.

Optionally, you can skip connecting port configuration by clicking **Next**.

Click the **Clear** button to remove the selected device from the configuration.

- Click the **Select Device** button to choose the device for which the port reservation will be associated. You should know the identity of the device to whose interface the new object will be associated before taking this step. For more information, see the section [Using the Device Selector](#).
  - After choosing the device, choose the **Interface** with which the port reservation will be bound. The drop-down list shows only interfaces that are most recently found to be available by Grid Manager during the last discovery cycle. This list will not include any ports that are Administratively Up and Operationally Up, or that are otherwise already assigned to other networks or objects.
  - The Wizard page also shows a list of any VLANs that are currently configured in the chosen device (**The following VLANs are configured**). This Wizard page allows only the assignment of an existing VLAN in the chosen device to the new port reservation.
  - Check the **Configure Port** checkbox to define specific port control settings for the port reservation.
  - Choose the **Data VLAN** and/or the **Voice VLAN** settings you may need for the port assignment. Depending on the selected device, you may or may not be able to apply VLAN settings.
  - Set the **Admin Status** to **Up** if you need to activate the port after assignment in the current task. All port control operations require CLI credentials to be entered into Grid Manager. Because some IPAM and DHCP objects will use port control features as part of object creation, CLI credentials are automatically leveraged as part of discovery and definition of port configurations such as Admin Up/Down status. Ensure you have the correct sets of CLI credentials for devices in your network.
  - Enter a **Description** for the port assignment. Infoblox recommends doing so to help other technicians to recognize the port assignment task.
8. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).

You need to assign a **Subscriber Member Site** to add subscriber service related extensible attribute in order to populate Subscriber Cache.

9. As the final step in the Add Fixed Address wizard, you define when Grid Manager creates the new object by scheduling it. You also schedule when the associated port configuration task executes.
- To create the new object and its associated port configuration immediately, select **Now**. The port control event is automatically synchronized to take place at the same time as the activation of the new object.
  - You can choose to have Grid Manager execute the port reservation task at the same time as the Fixed Address object creation. To do so, select **At same time as Host**.
  - You can choose to have Grid Manager execute the port reservation task at a later time by selecting **Later**. Choose a **Selected time** by entering or selecting a **Start Date** (click the calendar icon to choose a calendar date) and a **Start Time**, and choose a **Time Zone**.
10. Choose one of the following from the **Save&...** drop-down button menu:
- Click **Save & Close** to add the new object and close the wizard (this is the default).
  - Click **Save & Edit** to add the new object and launch the editor.
  - Click **Save & New** to add the new object and launch the wizard again to add another Fixed Address object.
11. Save the configuration and click **Restart** if it appears at the top of the screen.

#### **Note**

At any step during the wizard, you can click **Schedule for Later** to schedule the task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#). You cannot schedule this task when you are creating an object that is within a delegated scope.

For information on viewing IPv6 fixed addresses in a network, see [Viewing IPv4 DHCP Objects](#).

## Modifying IPv6 Fixed Addresses

To modify a fixed address:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *fixed\_address* checkbox, and then click the Edit icon.

2. The *Fixed Address* editor contains the following basic tabs from which you can modify data:
  - **General:** You can modify all the fields you filled out in the first step of the wizard described in [Adding IPv6 Fixed Addresses](#).
  - **Device Information:** You can change advisory Device Information settings for the object's port reservation; settings are described in the section [Configuring IPv6 Fixed Addresses](#).
  - **Discovery:** Checking the **Enable Discovery** checkbox informs NIOS to begin discovering the network after you click **Save and Close**. You manage discovery polling settings local to each fixed address from this page. For a complete overview of features on this page, see [Discovering Devices and Networks](#) and its subsections.
  - **Port Reservation:** Review and edit any device port reservations that may be defined for the current object, or create a new port reservation and schedule it. For a closer look, see the section [Port Control Features in Network Insight](#), and steps 5-9 in the section [Configuring IPv6 Fixed Addresses](#).
  - **IPv6 DHCP Options:** You can keep the inherited DHCP options or override them and enter unique settings for the fixed address. For information, see [Defining General IPv6 Properties](#).
  - **Filters:** You can keep the inherited IPv6 logic filters or override them and add a new IPv6 logic filter. For information, see [Applying Filters to DHCP Objects](#).
  - **Discovered Data:** You can view discovered data of this address, if any, in this tab. For information, see [Viewing Discovered Data](#).
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#).
3. Save the configuration and click **Restart** if it appears at the top of the screen.

or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Deleting IPv6 Fixed Addresses

To delete a fixed address, from the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* -> *fixed\_address* checkbox, and then click the Delete icon. You cannot delete multiple fixed addresses at the same time if the authority for one of the fixed addresses is delegated to a Cloud Platform Appliance.

## Configuring IPv6 Address Ranges

You can configure IPv6 ranges that are used to delegate IPv6 prefixes only, to assign IPv6 addresses only, or to delegate IPv6 prefixes and assign IP v6 addresses. When you define a DHCP range to delegate prefixes, the prefixes can be outside of the network where they are being defined. IPv6 ranges inherit their properties from their network, so each range in a subnet provides the same set of options to their DHCP clients.

Note that when an Infoblox DHCP server grants IPv4 leases, it starts from the last IP address in the range to the first. When the server grants IPv6 leases, it uses an algorithm based on the DUID of the client.

## Adding IPv6 Address Ranges

To add a an IPv6 address range:

1. Navigate to the IPv6 network to which you want to add an address range, and then select **Range** from the Add drop down menu.
 

or

From any panel in the DHCP tab, expand the Toolbar and click **Add** -> **Range** -> **IPv6**.
2. In the *Add IPv6 Range* wizard, select one of the following and click **Next**:
  - **Add IPv6 Range:** Select this to add an address range from scratch.
 

or
  - **Add IPv6 Range Using Template**

Click **Select Template** and select the template that you want to use. Note that when you use a template to create a DHCP range, the configurations of the template apply to the new range. The appliance automatically populates the address range properties in the wizard. You can then edit the pre-populated properties. For more information, see [About IPv6 Range Templates](#).

3. Complete the following:

- **Network:** Click **Select Network**. Grid Manager displays the network address here if you have only one network configured. When there are multiple networks, Grid Manager displays the *Select Network* dialog box from which you can select one.

Specify one of the following:

- **Address:** Select this if the address range is used to allocate IPv6 addresses only to DHCP clients, and then enter the start and end addresses in the range.
- **Prefix Delegated:** Select this if the DHCP server uses this address range to delegate IPv6 prefixes only to DHCP clients. Enter the start and end prefixes, and the prefix length.
- **Both:** Select this if the DHCP server delegates IPv6 prefixes and allocates IPv6 addresses from this range. Enter the start and end addresses in the range, and the start and end prefixes, and the prefix length.

Complete the following:

- **Name:** Enter a name for the address range.
- **Comment:** Enter additional information about the address range.
- **Disabled for DHCP:** Select this if you want to save the configuration for the address range but do not want to activate the address range yet. You can clear this checkbox when you are ready to allocate addresses from this range.

The **Cloud** section appears when the Cloud Network Automation license is installed on the Grid Master. For information, see [Deploying Cloud Network Automation](#). To delegate authority for this range, complete the following:

#### Delegate authority from the Grid Master

**DelegateTo:** This field indicates whether the authority for the range you want to create has already been delegated to a Cloud Platform Appliance. Click **Select** to choose the Cloud Platform Appliance to which you want to delegate authority. The *Member Selector* displays only Cloud Platform Appliances in the Grid. Click the member, and Grid Manager displays the member name next to this field. This cloud member now assumes authority for this range, and the Grid Master does not have authority any more. You can also click **Clear** to remove authority delegation from the selected Cloud Platform Appliance and return authority back to the Grid Master.

4. Click **Next** and select one of the following to provide DHCP services for the DHCP range:

- **None (Reserved Range):** Select this if you want to reserve this address range for static hosts. Addresses in this range cannot be allocated as dynamic addresses. You can allocate the next available IP from this range to a static host. This is selected by default.
- **Grid Member:** Select this if you want a Grid member to serve DHCP for this DHCP range. Select a Grid member from the drop-down list. The drop-down list displays only the Grid members that are associated with the network to which the DHCP range belongs.

5. (*Applies only to Network Insight*) Click **Next** to initiate or disable discovery of the new DHCP range.

- **Configuring one network:** Discovery settings include the following: **Enable Discovery** and **Immediate Discovery**, selecting a Probe member to perform the discovery; and **Polling Options**, which define how the network will be discovered by the Probe member using SNMP and CLI credentials. By default, all Polling Options discovery settings are inherited from the parent network unless you click **Override**. Polling Options govern the protocols used to query and collect information about the network devices being discovered. See the section [Configuring Discovery Properties](#) for a complete description of discovery Polling Options.

6. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).

7. Save the configuration and click **Restart** if it appears at the top of the screen.

or



Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Setting the Priority of IPv6 Address Ranges

The DHCP server allocates IP addresses from the configured DHCP ranges according to the order in which the ranges are listed. By default, ranges are listed according to their start addresses. You can move the ranges up and down in the list to change their order. For information about viewing DHCP ranges and other objects in a network, see [Viewing IPv6 DHCP Objects](#).

To change the order of DHCP ranges in a network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> *network*.
2. Expand the Toolbar and click **Order DHCP Ranges**.
3. In the *Order DHCP Ranges* dialog box, click the up and down arrows to move ranges up or down on the list. The Priority value changes accordingly. Click **OK** to save the configuration.

## Modifying IPv6 Address Ranges

To modify an IPv6 address range:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr\_range* checkbox, and then click the Edit icon.
2. The *DHCP Range* editor contains the following basic tabs from which you can modify data:
  - **General**: Modify the fields, except the network address, as described in [Adding IPv6 Address Ranges](#) above.  
When the Cloud Network Automation license is installed on the Grid Master, Grid Manager displays the following information in the **Cloud** section: **Cloud Usage**, **Owned By**, and **Delegated To**. You cannot modify these fields.
  - **Member Assignment**: Modify the Grid member that provides DHCP services for the DHCP range as described in [Adding IPv6 Address Ranges](#) above.
  - **Discovery**: You can enable and change discovery settings for the IPv6 range at any time after creating the range. Discovery settings include the following: **Enable Discovery** and **Immediate Discovery**, selecting a Probe member to perform the discovery; and **Polling Options**, which define how the network will be discovered by the Probe member using SNMP and CLI credentials. By default, all Polling Options discovery settings are inherited from the parent network unless you click **Override**. Polling Options govern the protocols used to query and collect information about the network devices being discovered. See the section [Configuring Discovery Properties](#) for a complete description of discovery Polling Options.
  - **Discovery Blackout**: Define extended time periods and regularly scheduled times when discovery and/or port configuration tasks will not take place on a network or DHCP range. Editing a network or range under DHCP, blackout settings apply only to the specified network or range. You also specify the scheduled time when the blackout period begins, and the duration of the blackout period. By default, the network inherits its discovery blackout settings from the Grid level. For related information, see [Defining Blackout Periods](#) and its subsections.
  - **IPv6 DHCP Options**: Keep active leases in a deleted DHCP range. For more information, see [Keeping Leases in Deleted IPv4 and IPv6 Networks and Ranges](#).
  - **Extensible Attributes**: You can add and delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions**: This tab appears only if you belong to a superuser admin group. For information, see [Managing Permissions](#).
3. Optionally, you can click **Toggle Advanced Mode** to display the following tabs from which you can modify advanced data.
  - **Exclusion Ranges**: Configure a range of IP addresses that the appliance does not use to assign to clients. You can use these exclusion addresses as static IP addresses. For more information, see [About Exclusion Ranges](#).
  - **Filters**: You can keep the inherited IPv6 logic filters or override them and add a new IPv6 logic filter. For information, see [Applying Filters to DHCP Objects](#).

Note that Grid Manager displays both the basic and advanced mode tabs the next time you log in to the GUI.

4. Save the configuration and click **Restart** if it appears at the top of the screen.  
or  
Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Deleting IPv6 Address Ranges

To delete an IPv6 address range:

- From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr\_range* checkbox, and then click the Delete icon.

## Viewing IPv6 DHCP Objects

You can view the DHCP objects in an IPv6 network by navigating to the **DHCP** tab -> **Networks** tab -> **Networks** panel, and then clicking the network link. This panel displays the following information about DHCP objects in the selected IPv6 network:

- **IP Address:** The IPv6 address of a DHCP object, such as a DHCP range, fixed address, or host configured for DHCP, or roaming host with an allocated IP address. For a DHCP range, this field displays the start and end addresses of the range. For a host that has multiple IP addresses, each IP address is displayed separately. Note that the appliance highlights all disabled DHCP objects in gray.
- **Type:** The DHCP object type, such as **IPv6 DHCP Range** or **IPv6 Fixed Address**.
- **Name:** The object name. For example, if the IP address belongs to a host record, this field displays the hostname.
- **Comment:** The information you entered for the object.
- **Site:** The site to which the DHCP object belongs. This is one of the predefined extensible attributes. You can select the following additional columns for display:
- **Priority:** Displays the priority of the DHCP range.
- **Disabled:** Indicates whether the network is disabled. You can also do the following in this panel:
- Modify some of the data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes. Note that some fields are read only.
- Sort the data in ascending or descending order by column.
- Create a bookmark for the object.
- Click **Go to IPAM View** to view information about the object in the **IPAM** tab.
- Delete or schedule the deletion of a selected object or multiple objects.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Print or export the data.

## Setting the Prefix Length Mode for DHCPv6

The prefix length mode determines the prefix selection rules employed by the DHCPv6 server when a DHCPv6 client sends an empty prefix with just a prefix length as a hint for the server to specify the required prefix length. This determines the prefix that gets allocated to the DHCPv6 client.

To set the prefix length mode:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.
2. In the editor, click the **General** tab -> **Advanced** tab.
3. In the IPv6 Properties area, from the **Prefix Length Mode** drop-down list, select from the following options:



- **Exact:** The DHCP server looks for a prefix with exactly the same length as the length requested by the client. If it cannot find a prefix that has the exact same length as the length requested, the server returns a status indicating that no prefixes are available. If you do not configure the prefix length mode, Exact is the default value.
- **Ignore:** The DHCP server ignores the length requested by the client and offers the first available prefix.
- **Prefer:** The DHCP server looks for prefixes that have the same length as the length requested by the client. If it does not find a prefix of the same length, it offers the first available prefix of any length.
- **Minimum:** The DHCP server looks for a prefix that has the same length as the requested length. If it does not find such a prefix, it returns a prefix whose length is greater than (that is, longer than) the requested value. If it does not even find a prefix of a greater length, it returns a status indicating that no are prefixes available. For example, if the client requests a prefix length of /60, and the DHCP server has prefixes of lengths /56 and /64 available, it offers a prefix of length /64.
- **Maximum:** The DHCP server looks for a prefix that has the same length as the requested length. If it does not find such a prefix, it returns a prefix whose length is less than (that is, shorter than) the requested value. If it does not even find a prefix of a shorter length, it returns a status indicating that no prefixes are available. For example, if the client requests a length of /60, and the server has prefixes of lengths /56 and /64 available, it offers a prefix of length /56.

4. Click **Save & Close**.

## Configuring DHCP Failover

This section explains how to configure DHCP failover associations. It contains the following topics:

- [DHCP Failover](#)
- [Configuring Failover Associations](#)
- [Managing Failover Associations](#)

### DHCP Failover

You can create a failover association between two DHCP servers (a primary server and a secondary) and assign the failover association to serve an IPv4 DHCP range. When you set up a failover association, you greatly reduce DHCP service downtime if one of your DHCP servers is out of service. You can better manage IP address requests by making two servers available for DHCP services. You can also configure one of the servers to assume full DHCP services when you know the other server may go out of service for a period of time.

You can configure two NIOS appliances, or one appliance and one external server, to form a failover association. The pairing of a primary and secondary server is called a peer association. The failover peers establish a TCP connection for their communication. They share a pool of IP addresses that they allocate to hosts on their networks based on load balancing. Load balancing is a technique to split the address allocation workload evenly across the two DHCP servers. You can assign a DHCP failover association to serve DHCP ranges in a network. A DHCP failover association can serve DHCP ranges that belong to one network view only. It cannot serve ranges in different network views.

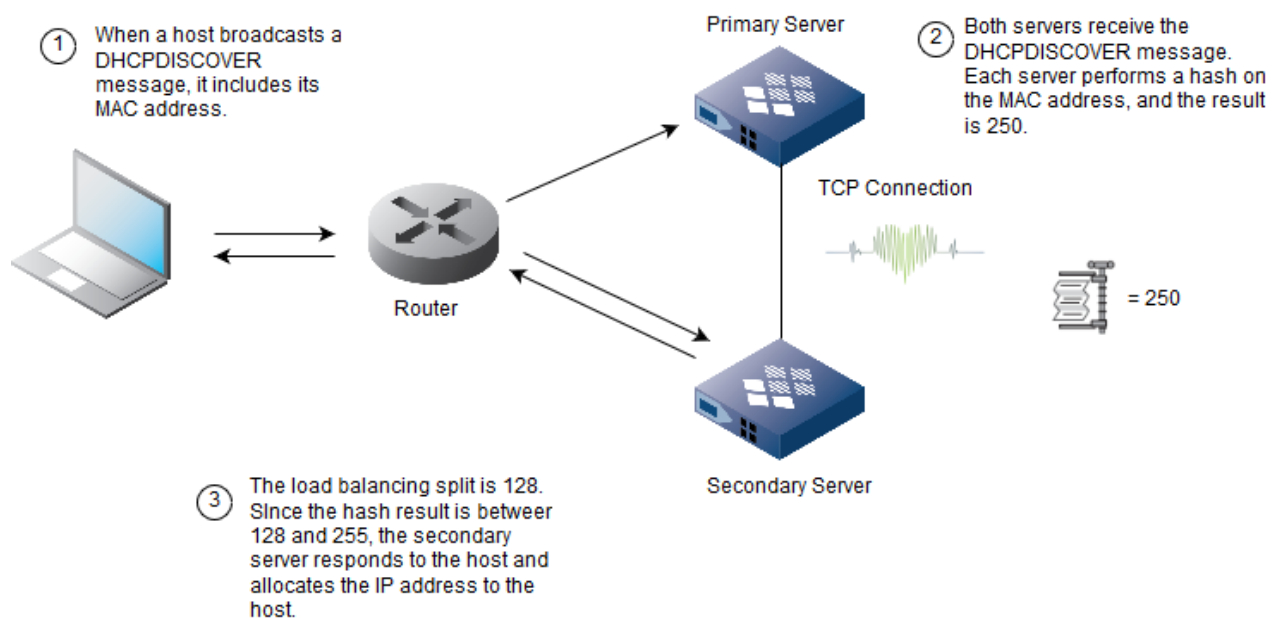
#### Note

When you assign a failover association to serve DHCP ranges and networks, NIOS denies dynamic BOOTP clients by default, regardless of whether you select or deselect the **Deny BOOTP Requests** option from Grid Manager. However, if the DHCP ranges or networks are assigned to a single DHCP server (not a failover association), NIOS does not automatically deny dynamic BOOTP clients. In this case, you must manually select the **Deny BOOTP Requests** option through Grid Manager to ensure that NIOS denies BOOTP requests to avoid problems such as receiving two IP addresses for the same network device. For information about how to deny BOOTP requests, see [Configuring IPv4 BOOTP and PXE Properties](#).

## Failover Association Operations

When a host broadcasts a DHCPDISCOVER message, it includes its MAC address. Both the primary and secondary peers receive this message. To determine which server should allocate an IP address to the host, they each extract the MAC address from the DHCPDISCOVER message and perform a hash operation. Each server then compares the result of its hash operation with the configured load balancing split. The split is set to 50% by default to ensure an even split between the two servers. When the split is 50%, the primary server allocates the IP address if the hash result is between 1 and 127, and the secondary server allocates the IP address if the hash result is between 128 and 255. As a server allocates an IP address, it updates its peer so their databases remain synchronized. As shown in the following figure *Load Balancing and IP Addresses Allocation*, when a host broadcasts a DHCPDISCOVER message, both the primary and secondary servers receive the message. They perform a hash operation on the MAC address in the DHCPDISCOVER message, and the result is 250. Since the load balancing split is 50% and the hash result is 250, the secondary server responds to the host with a DHCP OFFER message. The secondary peer allocates an IP address from its assigned pool of IP addresses. It then sends a lease update message to the primary server so that the primary server knows how the address is assigned and can properly take over if the secondary server fails.

### *Load Balancing and IP Addresses Allocation*



Related topic

[Configuring DHCP Failover](#)

## Configuring Failover Associations

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.
2. Click **Toggle Advanced Mode** if the editor is in basic mode. When the additional tabs appear, select the **General Advanced** tab to complete the following:
  - **Failover Port:** You can modify the port number that members use for failover associations. You can use any available port from 519 to 647. The default is 647 for a new installation and 519 for an upgrade.

The following are tasks and guidelines for configuring a DHCP failover:

1. Identify the primary and secondary DHCP servers and ensure that the appliances are set up correctly for the failover association, using the following guidelines:
  - Configure a failover association using two NIOS appliances, or a NIOS appliance and an ISC DHCP compliant server.
  - One of the DHCP servers must be an independent appliance or in an Infoblox Grid.
  - The DHCP servers do not have to be in the same geographic location.
  - The clocks on both servers must be synchronized. This happens automatically when both servers are on the same Grid.
  - Both servers must use the same version of the DHCP configuration file. This happens automatically when both servers are on the same Grid.
  - If you use firewalls on your networks, ensure that the firewalls allow TCP port 647 between the servers, and that TCP port 7911 is open for partner down operations.
  - Each pair of DHCP servers can participate in only one failover association. An appliance can participate in more than one failover association, as long as it is with a different peer.

Configure the same DHCP properties on the primary and secondary servers, as described in [Configuring General IPv4 DHCP Properties](#).

- Both the primary and secondary servers must have the same operational parameters, and they must be able to receive DHCPDISCOVER messages that hosts broadcast on the networks.
- If you change any of the DHCP failover parameters for a peer association definition, you must make the same changes on both the primary and secondary servers.

Note that if both the primary and secondary servers are in a Grid, you configure the properties on the failover association and the configuration applies to both servers.

2. Create a failover association and configure load balancing between the servers. For information, see [Adding Failover Associations](#) below.
  - Ensure that you use the same failover association name on both the primary and secondary servers.
  - The appliance assigns default values to the failover timers and triggers. In general, these default values serve the purpose of a failover. Do not change these values unless you understand the ramification of the changes. For example, when one of the peers in a failover association fails, the other peer goes into a COMMUNICATIONS-INTERRUPTED state, and the lease time changes to the MCLT (Maximum Client Lead Time). You should consider how the MCLT affects the lease time when a failover occurs if you want to change this value.
3. Assign the failover association to the DHCP ranges in the same network view. Failover associations can serve only IPv4 DHCP ranges. For information, see [Configuring IPv4 Address Ranges](#).
  - If you configure a shared network, and the subnets in the shared network contain ranges served by a DHCP failover association, both the primary and secondary DHCP server must have the same shared networks defined, containing the same networks and DHCP ranges.  
If you have multiple networks that are in a shared network and you plan to use a DHCP failover, you must use the same failover association and specify the same peers on all the networks in the shared network.
4. Enable DHCP on the primary and secondary servers AFTER you complete all the configurations. For information, see [Managing Failover Associations](#).



#### Note

When you set up a failover association for the first time, ensure that both servers are up and running and their databases are synchronized before they can start assigning IP addresses.

When you configure a failover association, the appliance assigns default values for timers and triggers, such as the MCLT and the maximum number of "unacked" packets. A failover may occur when some of the timers expire or when a failover peer goes out of service. When a failover occurs, the functional peer takes over and assigns IP addresses with the lease time set to the MCLT. When the server that is offline comes back online, it synchronizes its database with its peer before it starts allocating IP addresses.

## Adding Failover Associations

To add a DHCP failover association, perform the following procedures on both the primary and secondary servers:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **IPv4 Failover Associations** section, and then click the Add icon.  
or  
Expand the Toolbar and click **Add -> IPv4 Failover Association**.
2. In the *Add Failover Association* wizard, complete the following:
  - **Name:** Enter a unique name for the failover association. The failover association name is case sensitive. Enter the same name on both the primary and secondary servers. The appliance validates the names on both servers. The names must be exactly the same. If they do not match, the failover association goes into disconnect mode.
  - **DHCP Failover Primary:** Select one of the following. The default is **Grid Member**.
    - **Grid Member:** Click **Select member**. In the *Select Member* dialog box, select the primary server and click the Select icon.
    - **MS Server:** Click **Select Server**. In the *Microsoft Server Selector* dialog box, select the Microsoft server that supports DHCP failover. Note that only certain versions of Microsoft servers support failover. This dialog box also displays the following columns: **IP address**, **Comment**, and **Site**. The failover association requires at least two Microsoft servers. NIOS displays an error message when you select **MS Server** if there is only one Microsoft server. When you add or modify a failover association through Grid Manager, the appliance displays a message to ensure the connectivity between two Microsoft servers. If you want a failover association to be in a normal state, you must ensure an end-to-end connectivity. If the connectivity fails, you must re-establish connectivity between the two Microsoft servers and the failover association follows the recovery states based on the DHCP failover protocol.  
You can configure Microsoft primary and secondary servers using this wizard only. You cannot edit Microsoft primary and secondary servers after you configure them. You must delete and reconfigure the primary or secondary Microsoft server for a failover association.
  - **External Server IP Address:** Select this to use an external ISC DHCP compliant server as the primary server. Enter the IP address of the primary server in the field.
  - **DHCP Failover Secondary:** Select one of the of following. The default is **Grid Member**.
    - **Grid Member:** Click **Select member**. In the *Select Member* dialog box, select the secondary server and click the Select icon.
    - **MS Server:** NIOS selects this automatically when you set the **DHCP Failover Primary** to **MS Server**. Click **Select Server**. In the *Microsoft Server Selector* dialog box, select the Microsoft server that supports DHCP failover. Note that only Microsoft Windows Server 2012 or later versions support synchronization of failover relationships. This dialog box also displays the following columns: **IP address**, **Comment**, and **Site**.  
The primary and secondary Microsoft servers that you select in a failover relationship must be in the same network view. For more information, see [About Microsoft DHCP Failover Relationships](#).
    - **External Server IP Address:** Select this to use an external ISC DHCP compliant server as the secondary server. Enter the IP address of the secondary server in the field.  
You cannot select **External Server IP Address** for both the primary and secondary servers. One of the servers must be an independent appliance or in an Infoblox Grid.
  - **Comment:** Enter useful information about the failover association.
3. Click **Next** and do the following to control the IP address allocation between the peers and how they switch from one to the other based on the configuration:
  - **Failover mode:** This is valid for Microsoft Management only. Select a failover mode. You can either select **Hot standby** or **Load balancing**. When you select **Hot standby**, the secondary server is set to **Standby** by default and the slider move to the position at 95%. The slider moves to 50% when you set the failover mode to **Load balancing mode**. You can synchronize and manage a failover mode that is operating in **Hot standby** mode. The primary partner is the active server that first creates the relationship if you use the **Load balancing** or the **Hot standby** mode.  
When you configure a failover association, the slider changes its position based on the Failover mode you select. When you edit failover settings, the slider remains in the Balanced position, at 50%, by default. For more information about modifying failover associations, see [Modifying Failover Associations](#).

- **Load Balancing Data:** Adjust the slider to determine which server should handle more IP address requests. The default is 50%. When you move the slider all the way to the left, the primary server responds to all IP address requests. The opposite applies when you move the slider all the way to the right. Infoblox recommends that you use the default (50/50) to enable the primary and secondary servers to respond to IP address requests on an equal basis.
  - **Lease Deletion:** Select the following to override settings at the Grid and member levels.
    - **Keep leases from deleted ranges until one week after expiration:** When you select this and delete a DHCP range with active leases, the appliance stores these leases up to one week after they expire. When you add a new DHCP range that includes the IP addresses of these active leases, the appliance automatically restores the leases.
  - **Secondary role:** This is valid for Microsoft Management only. Note that the secondary role is available only in **Hot standby** mode. The appliance displays **Standby** by default and you cannot edit this value.
  - **Maximum client lead time:** Specify the maximum client lead time in minutes or hours. The default is one hour. Select **Minutes** or **Hours** from the drop-down list. This specifies the maximum amount of time the server waits before assuming control.
  - **Enable switchover interval:** This is valid for Microsoft Management only. Select this to automatically change the state to partner down after a specified period. NIOS does not support the "partner down" state for Microsoft DHCP failover association.
  - **State switchover interval (Minutes):** This is valid for Microsoft Management only. Specify the amount of time after which the server must change the state. The default is 60 minutes.
  - **Enable Authentication:** This is valid for Microsoft Management only. Select this if you want to secure the communication between failover partners.
  - **Shared Secret:** This is valid for Microsoft Management only. Enter a shared secret that can be used to authenticate the communication between failover partners. You can specify a shared secret only if you enable authentication.
4. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).
  5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Best Practices for Restoring Microsoft DHCP Server Data

When you synchronize two Microsoft servers in read/write mode for a Microsoft DHCP failover association, the appliance might override new data. For example, assume that you are synchronizing data from one Microsoft server, which has latest data. But after the restore operation, the appliance will synchronize data from a Microsoft server that has old data. If you synchronize old data, it might override new data. Infoblox recommends that you follow the steps mentioned below to prevent new data being replaced with the old one:

1. Disable DHCP synchronization for the Microsoft server before you restore data to the Microsoft server.
2. Restore the Microsoft DHCP server.
3. From the Microsoft server, which has the latest data, replicate its DHCP failover association and scopes to the restored Microsoft server. This ensures that both Microsoft servers have the same latest data.
4. Re-enable DHCP synchronization for the restored Microsoft server. NIOS will resynchronize with the Microsoft server.

## Managing Failover Associations

After you establish a failover association, you can monitor its status periodically to ensure that it is functioning properly. You can also delete a failover association when it is not assigned to any DHCP range.

See the following sections on how to manage failover associations:

- [Modifying Failover Associations](#)
- [Monitoring Failover Associations](#)
- [Deleting Failover Associations](#)
- [Setting a Peer in the Partner-Down State](#)
- [Performing a Force Recovery](#)
- [Recovering DHCP Failover Associations](#)

Under special circumstances, you can manually adjust the configuration of a failover association. For example, when you know in advance that a peer will be out of service for an extended period of time, you can manually set the functional peer in a PARTNER-DOWN mode. This allows the functional partner to assume all leases and be able to allocate addresses to client requests in full capacity. In addition, when you suspect the databases in a failover association are not synchronized, you can consider doing a force recovery (after you consult with Infoblox Technical Support or your Infoblox representative) so the secondary server can completely rebuild its lease table with updates from the primary server. See the following sections on how to set a peer to the partner-down mode and perform a force recovery:

- [Setting a Peer in the Partner-Down State](#)
- [Performing a Force Recovery](#)

## Modifying Failover Associations

To modify a failover association:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Failover Associations** -> *failover\_association* checkbox, and then click the Edit icon.
2. The *DHCP Failover Association* editor contains the following tabs from which you can modify data:
  - **General:** In the **Basic** tab, modify the fields as described in [Adding Failover Associations](#). In the **Advanced** tab, complete the following to modify the port number you use for the failover association:
    - **Failover Port:** Click **Override** to enter a port number for the failover association. You can use any available port from 1 to 63999. The default is 647 for a new installation and 519 for an upgrade.
  - **Triggers:** Before editing the triggers and timers, ensure that you understand the ramification of the changes. Improper configuration of the triggers can cause the failover association to fail. For information about the fields in the **Basic** tab, see [Adding Failover Associations](#). The following are the triggers in the **Advanced** tab:
    - **Max Response Delay Before Failover(s):** Specifies the maximum duration of time (in seconds) before a failover enters the Communications-Interrupted state after failing to hear from its peers. The duration must be long enough to prevent frequent connections and disconnections from the DHCP failover peers, yet short enough so that the transient network failure will not keep the peers out of contact for an extended duration. The recommended default is 60 seconds.
    - **Max Number of Unacked Updates:** Specifies the number of "unacked" packets the server can send before a failover occurs. The default is 10 messages.
    - **Max Client Lead Time (s):** Specifies the length of time that a failover peer can renew a lease without contacting its peer. The larger the number, the longer it takes for the peer to recover IP addresses after moving to the **PARTNER-DOWN** state. The smaller the number, the more load your servers experience when they are not communicating. The default is 3600 seconds.
    - **Max Load Balancing Delay (s):** Specifies the cutoff after load balancing is disabled. The cutoff is based on the number of seconds since a client sent its first **DHCPDISCOVER** message. For instance, if one of the failover peers gets into a state where it is busy responding to failover messages but is not responding to other client requests, the other peer responds to the client requests when the clients retry. This does not cause a failover. The default is three seconds.
  - **Failover Settings:** This is valid for Microsoft Management only. Modify failover association settings. For information, see [Configuring Failover Associations](#). If you modify failover settings from secondary Microsoft server settings, the appliance does not update failover settings on NIOS for the following reasons:
    - When DHCP synchronization is disabled for primary Microsoft server, you must enable DHCP synchronization for primary Microsoft server to reflect the settings on NIOS.
    - The primary synchronization interval must be completed. For example, consider that you are modifying failover settings from secondary Microsoft server settings where the synchronization interval for primary server is five minutes, and the time interval for the secondary server is one minute. In this case, failover settings are updated on NIOS only after the primary server synchronization interval, which is five minutes.
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with a failover association. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).



## Monitoring Failover Associations

After you configure a failover association, the peers establish a TCP connection for communication. In a normal operational state, they send keepalive messages and database updates every time they grant a lease. However, there are times when the failover association experiences problems and goes into a state other than **NORMAL**. You can monitor the overall state of a failover association and the individual status of the peers to verify that the servers are operating and communicating properly.

Both peers in a failover association maintain the same DHCP fingerprinting state (enabled or disabled) even when one of the peers fails or becomes operational again. Note that both peers must be in the same Grid for the fingerprinting state to stay the same. For information about DHCP fingerprinting, see [About DHCP Fingerprints](#).

In this panel, you can also modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).

To monitor the failover association status:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **IPv4 Failover Associations** section. Grid Manager displays the list of failover associations and their overall status.
2. To view detailed information about a failover association, select the *failover\_association* checkbox, and then click the Show Status icon.
3. In the *Failover Association Status* dialog box, Grid Manager displays the overall status of the failover association and the status of both the primary and secondary servers.

The failover association can be in one of the following states:

- **OK** (green): The failover association is functioning properly.
- **DEGRADED** (yellow): The failover association is degraded when one of the peers is giving out limited addresses.
- **FAILURE** (red): The failover association is not functioning, may be because it is not completely configured. The peers are not assigning IP addresses.  
For each peer, Grid Manager displays the hostname or IP address, the status, and event date. The peer can be in one of the following states:
  - **STARTUP**: The server is starting up.
  - **NORMAL**: The server is in a normal operational state in which it responds to its load balancing subset of DHCP clients.
  - **PAUSED**: This state allows a peer to inform the other peer that it is going out of service for a short period of time so the other peer can immediately transition to the **COMMUNICATIONS-INTERRUPTED** state and start providing DHCP service to DHCP clients.
  - **COMMUNICATIONS-INTERRUPTED**: The servers are not communicating with each other. Both servers provide DHCP service to DHCP clients from which they receive DHCP requests.
  - **PARTNER-DOWN**: The server assumes control of the DHCP service because its peer is out of service.
  - **RECOVER**: The server is starting up and trying to get a complete update from its peer and discovers that its peer is in the **PARTNER-DOWN** state.
  - **RECOVER-WAIT**: The server has got a complete update from its peer and is waiting for MCLT period to pass before transitioning to the **RECOVER-DONE** state.
  - **RECOVER-DONE**: The server completed an update from its peer.
  - **POTENTIAL-CONFLICT**: The peers are not synchronized due to an administrative error or an incorrect state transition. Check the failover configuration and correct the error.
  - **CONFLICT-DONE**: This is a temporary state that the primary server enters after it received updates from the secondary server when it was in the **POTENTIAL-CONFLICT** state.
  - **RESOLUTION-INTERRUPTED**: The server responds to DHCP clients in a limited way when it is in this state.
  - **UNKNOWN**: The DHCP server is in an unknown state. The failover association is not functioning properly, may be because it is configured improperly. For example, failover association is not assigned to any DHCP range.
  - **SHUTDOWN**: This state allows a peer to inform the other peer that it is going out of service for a long period of time so the other peer can immediately transition to the **PARTNER-DOWN** state and completely assume control of the DHCP service.





NIOS does not support **PARTNER-DOWN** and Force Recovery for a Microsoft DHCP failover association.

## Deleting Failover Associations

You cannot delete a failover association if it is currently assigned to a DHCP range. If you want to delete a failover association, ensure that it is not assigned to any DHCP range.

To delete a failover association:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Failover Associations** -> *failover\_association* checkbox, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.  
The appliance puts the failover association in the recycle bin, if enabled.

## Setting a Peer in the Partner-Down State

If one of the peers in a failover association is out of service for an extended period of time, you should consider putting the functional peer in the **PARTNER-DOWN** state. When you place the functional peer in the **PARTNER-DOWN** state, it assumes full DHCP services for the networks. Since the functional server may not receive all the updates from its peer, it extends all the leases on the MCLT. Once the following conditions are met, the functional peer provides DHCP services autonomously:

- It has reclaimed all the leases that belonged to its peer.
- The MCLT has passed.

When the peer that is offline comes back online, it synchronizes with the functional peer and reestablishes the communication before it provides DHCP services to the clients.



### Warning

*Before you put a peer in the **partner-down** state, ensure that the other peer is indeed out of service. If both the primary and secondary servers are operational when you place one of them in the partner-down mode, both servers may stop issuing leases for a minimum of time defined in the MCLT.*

To set a peer in the **PARTNER-DOWN** state:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Failover Associations** -> *failover\_association* checkbox.
2. Expand the Toolbar and click **Set Partner Down**.
3. In the *Set Failover Association Partner Down* dialog box, select one of the following:
  - **Primary**: Select this if the secondary server is out of service.
  - **Secondary**: Select this if the primary server is out of service.
4. Click **OK**.



### Note

You cannot place the functional peer in the **PARTNER-DOWN** state for a Microsoft DHCP failover in NIOS.

## Performing a Force Recovery

When the primary and secondary peers are not synchronized, you can perform a force recovery to set the primary server in the **PARTNER-DOWN** state while putting the secondary server in the **RECOVER** state. During a force recovery, all leases in the databases are resynchronized. When you perform a force recovery, the secondary server does not serve any DHCP leases for a minimum of the MCLT while it resynchronizes with the primary server. Before you perform a force recovery, consult with Infoblox Technical Support or your Infoblox representative to ensure that the force recovery is appropriate for the situation.

To perform a force recovery:

1. From the **Data Management** tab, select the **DHCP** tab-> **Members** tab-> **Failover Associations** -> *failover\_association* checkbox.
2. Expand the Toolbar and click **Force Recovery State**.
3. In the *Force Secondary Peer Recovery State* dialog box, click **OK**.

The appliance synchronizes the databases on the primary and secondary servers.



**Note**

You cannot place the functional peer in the **PARTNER-DOWN** state for a force recovery in NIOS.

## Recovering DHCP Failover Associations

During a conflict resolution, when the primary peer of the DHCP failover association is in the **CONFLICT-DONE** state and the secondary peer in the **POTENTIAL-CONFLICT** state, the secondary peer might experience problems (such as restarting, network outage, etc.) and goes into an invalid state. This results in a deadlock state for the failover association, causing a DHCP service outage. When the failover association is in a deadlock state, you can perform a recovery for the failover association. You can run the recovery for one failover association at a time and when the primary member is in the **CONFLICT-DONE** state. This feature is supported for Infoblox appliances only and not for any other external DHCP servers.



**Note**

When the failover recovery is in progress, the DHCP service on both peers are disabled and you cannot enable the DHCP service until the failover recovery is successfully completed. You can view the logs of the failover recovery process in the syslog and infoblox.log file.

To recover a DHCP failover association:

1. From the **Data Management** tab, select the **DHCP** tab-> **Members** tab -> **Failover Associations** -> *failover\_association* checkbox.
2. Expand the Toolbar and click **Recovery from Deadlock State**.
3. In the *Failover Recovery Progress* dialog box, click **Start** to start the recovery of the failover association from the deadlock state.
4. In the confirmation dialog box, click **Yes**.



**Note**

After you start the failover recovery, you cannot revert the changes.

Grid Manager starts the failover recovery and you can view the following information in the *Failover Recovery Progress* dialog box:

- **Failover association:** The name of the failover association.
- **Primary:** The hostname or IP address of the primary server.
- **Secondary:** The hostname or IP address of the secondary server.
- **Number of leases to be processed:** The total number of leases to be processed.
- **Number of leases processed:** The number of leases that have been processed.
- **Current Status:** Displays the current status of the failover recovery process. The current status can be one of the following:
  - **Pending:** The failover recovery is initiated for a failover association and the recovery process will start soon.
  - **Calculating:** The appliance calculates the total amount of leases to be processed.
  - **Applying:** The appliance looks for conflicts and tries to resolve the conflicts.
  - **Completed:** The failover recovery is completed successfully.
  - **Failed:** The failover recovery fails.

Grid Manager also displays the reason for the failure if that happens.

After successful completion of the failover recovery, you must restart both the primary and secondary peers to bring them back to the **CONFLICT-DONE** state.

You can stop the failover recovery operation by clicking **Stop** in the *Failover Recovery Progress* dialog box before the recovery process is complete.



#### Note

If for any reasons the recovery is blocked when the operation is in progress, you can cancel the current operation and start the recovery for the failover association again.

## Configuring DHCP Filters

To control how the appliance allocates IPv4 or IPv6 addresses, you can define DHCP filters and apply them to Grid, members, DHCP ranges, range templates, IPv4 network containers, IPv4 networks, shared networks, IPv4 / IPv6 network templates, IPv4 / IPv6 fixed addresses, IPv4 / IPv6 fixed address templates, IPv4 / IPv6 reservations, IPv4 / IPv6 reservation templates, and IPv4 / IPv6 host addresses. You can override filters set at an upper level and apply a new logic filter. Depending on your configuration, DHCP filters screen requesting clients by matching MAC addresses, relay agent identifiers, DHCP options, or DHCP fingerprints you define in the filters. If you configure DHCP servers in the Grid to send authentication requests to a RADIUS authentication server group, you can also filter requests by matching the authentication results. (For information about this feature, see [Authenticated DHCP](#).)

When you define DHCP filters, you classify DHCP clients based on the information provided by the clients or by the RADIUS server. When you apply filters to an address range, the appliance responds to your address requests based on your configuration. The appliance also decides which DHCP options to return to the matching clients based on how you apply the filters. For more information, see [Applying Filters to DHCP Objects](#).

You can use filters to control address allocation based on your network requirements. For example, you can use DHCP filters to screen unmanaged hosts on a network by denying their address and option requests. If you have multiple DHCP address ranges on the same network and you want to assign IP addresses from specific address ranges to specific clients, you can use filters to screen the address assignments. For information, see [IP Address Allocation](#).

The appliance supports the following filters:

- **MAC address filters:** Use MAC addresses as matching criteria for granting or denying address requests. For information, see [Configuring MAC Address Filters](#).
- **Relay agent filters:** Identify remote hosts by matching the relay agent identifiers in the DHCPDISCOVER messages. For information, see [About Relay Agent Filters](#).
- **Option filters:** Classify hosts by matching the DHCP options and values sent by the requesting hosts. For information, see [Configuring Option Filters](#).
- **DHCP fingerprint filter:** Identify remote clients by matching the option number sequence or vendor ID sent in option 55 and 60 of the DHCP request against the DHCP fingerprints cached on the system. For information about DHCP fingerprint filters, see [Configuring DHCP Fingerprint Filters](#). For information about DHCP fingerprint detection, see [DHCP Fingerprint Detection](#).
- **NAC filters:** Use authentication results from a RADIUS authentication server group as matching criteria for granting or denying address requests. For information, see [Authenticated DHCP](#).

You can use MAC, option, and NAC filters to define DHCP options that matching clients can receive. Depending on how you apply a filter, all DHCP clients with matching criteria can receive all or some of the DHCP options defined in the filter. DHCP options defined for a matching filter supersede those defined at the Grid, member, network, and DHCP range levels. Options defined for a filter that is in the Class Filter List of an address range supersede those defined in the Logic Filter List. For more information about how the appliance returns options and how to apply DHCP filters, see [Applying Filters to DHCP Objects](#).

This section explains how to configure IPv4 DHCP filters. It contains the following topics:

- [IP Address Allocation](#)
- [IP Address Allocation Using Filters](#)
- [Configuring MAC Address Filters](#)
- [About Relay Agent Filters](#)
- [Configuring Option Filters](#)
- [Configuring DHCP Fingerprint Filters](#)

- [Applying Filters to DHCP Objects](#)
- [Managing DHCP Filters](#)

## IP Address Allocation

When a DHCP client requests an IP address, the NIOS appliance draws an address from an address range associated with the network segment for that client. Because you define that range, you can thereby control the IP address (within the defined range) and the associated TCP/IP settings that the client receives.

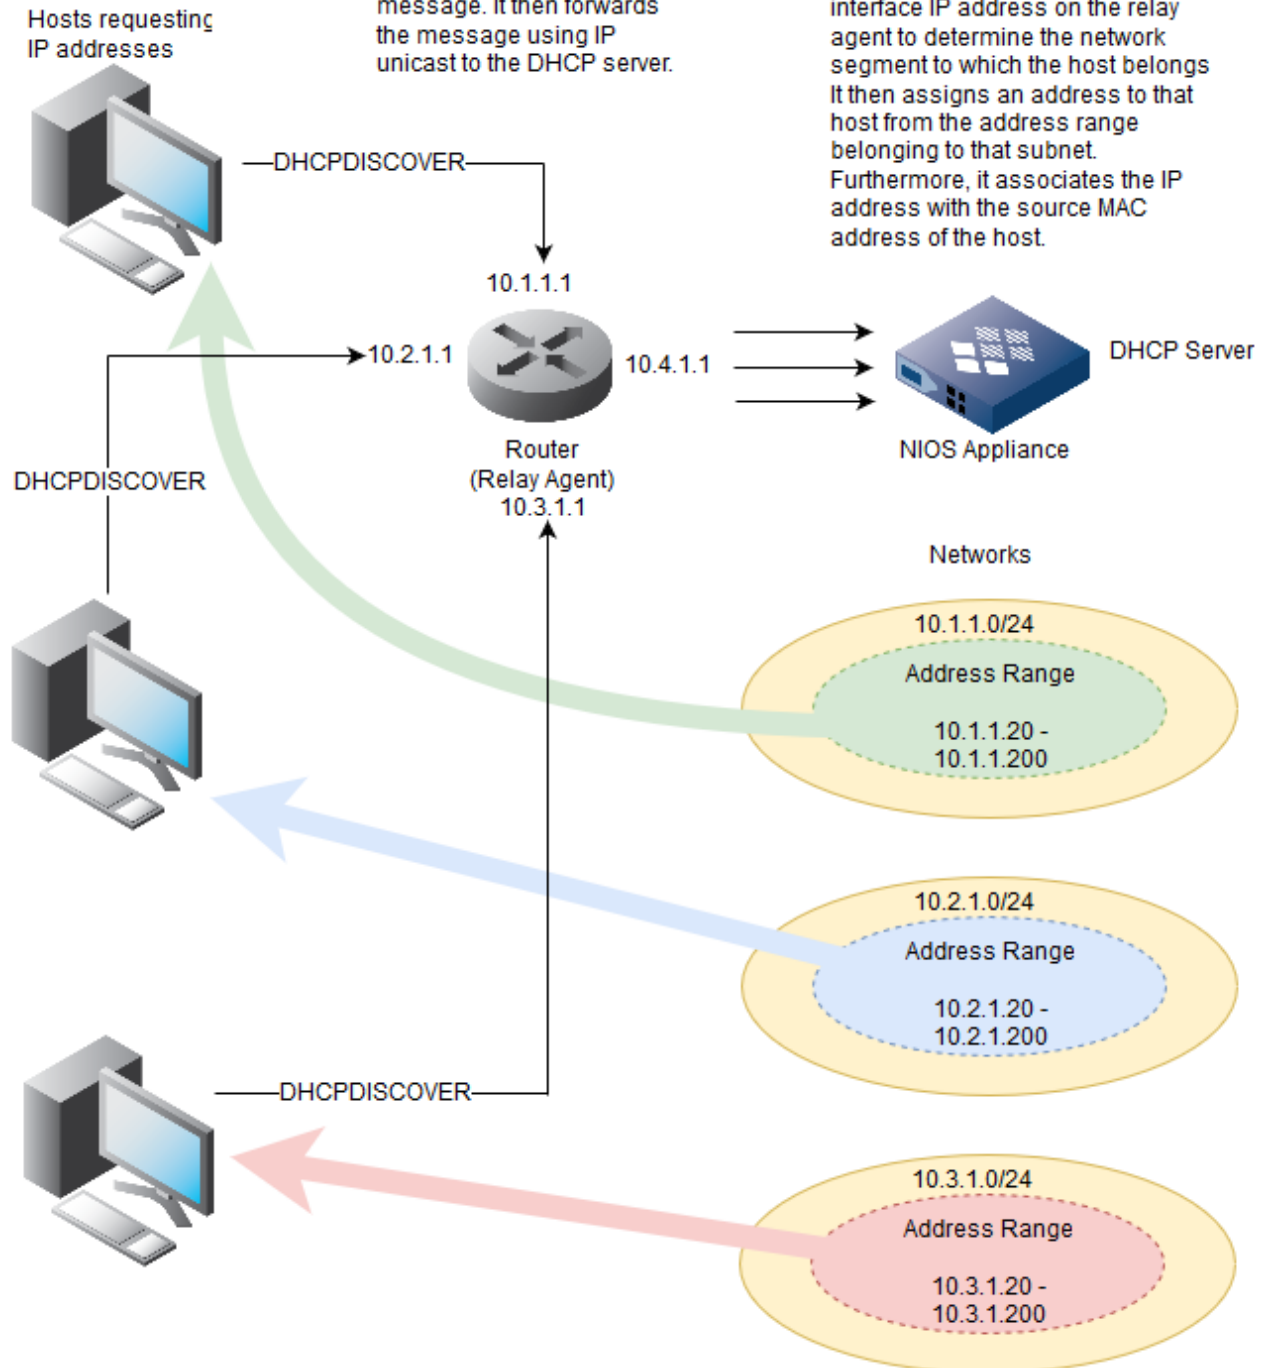
In the following figure [Requesting Addresses – DHCPDISCOVER Messages](#), three hosts — each in a different subnet — request an IP address. Each one broadcasts a DHCPDISCOVER message, which includes its MAC address. When the router, which also functions as a DHCP relay agent, receives the message, it adds the IP address of the interface on which the message arrives and forwards the message to the DHCP server — or servers — previously configured on the router. When the NIOS appliance receives the message, it uses the ingress interface IP address of the router to determine the network segment to which the host belongs and associates the MAC address of the requesting host with an IP address from an address range for that network.

*Requesting Addresses – DHCPDISCOVER Messages*

① When each host broadcasts a DHCPDISCOVER message, it includes its MAC address.

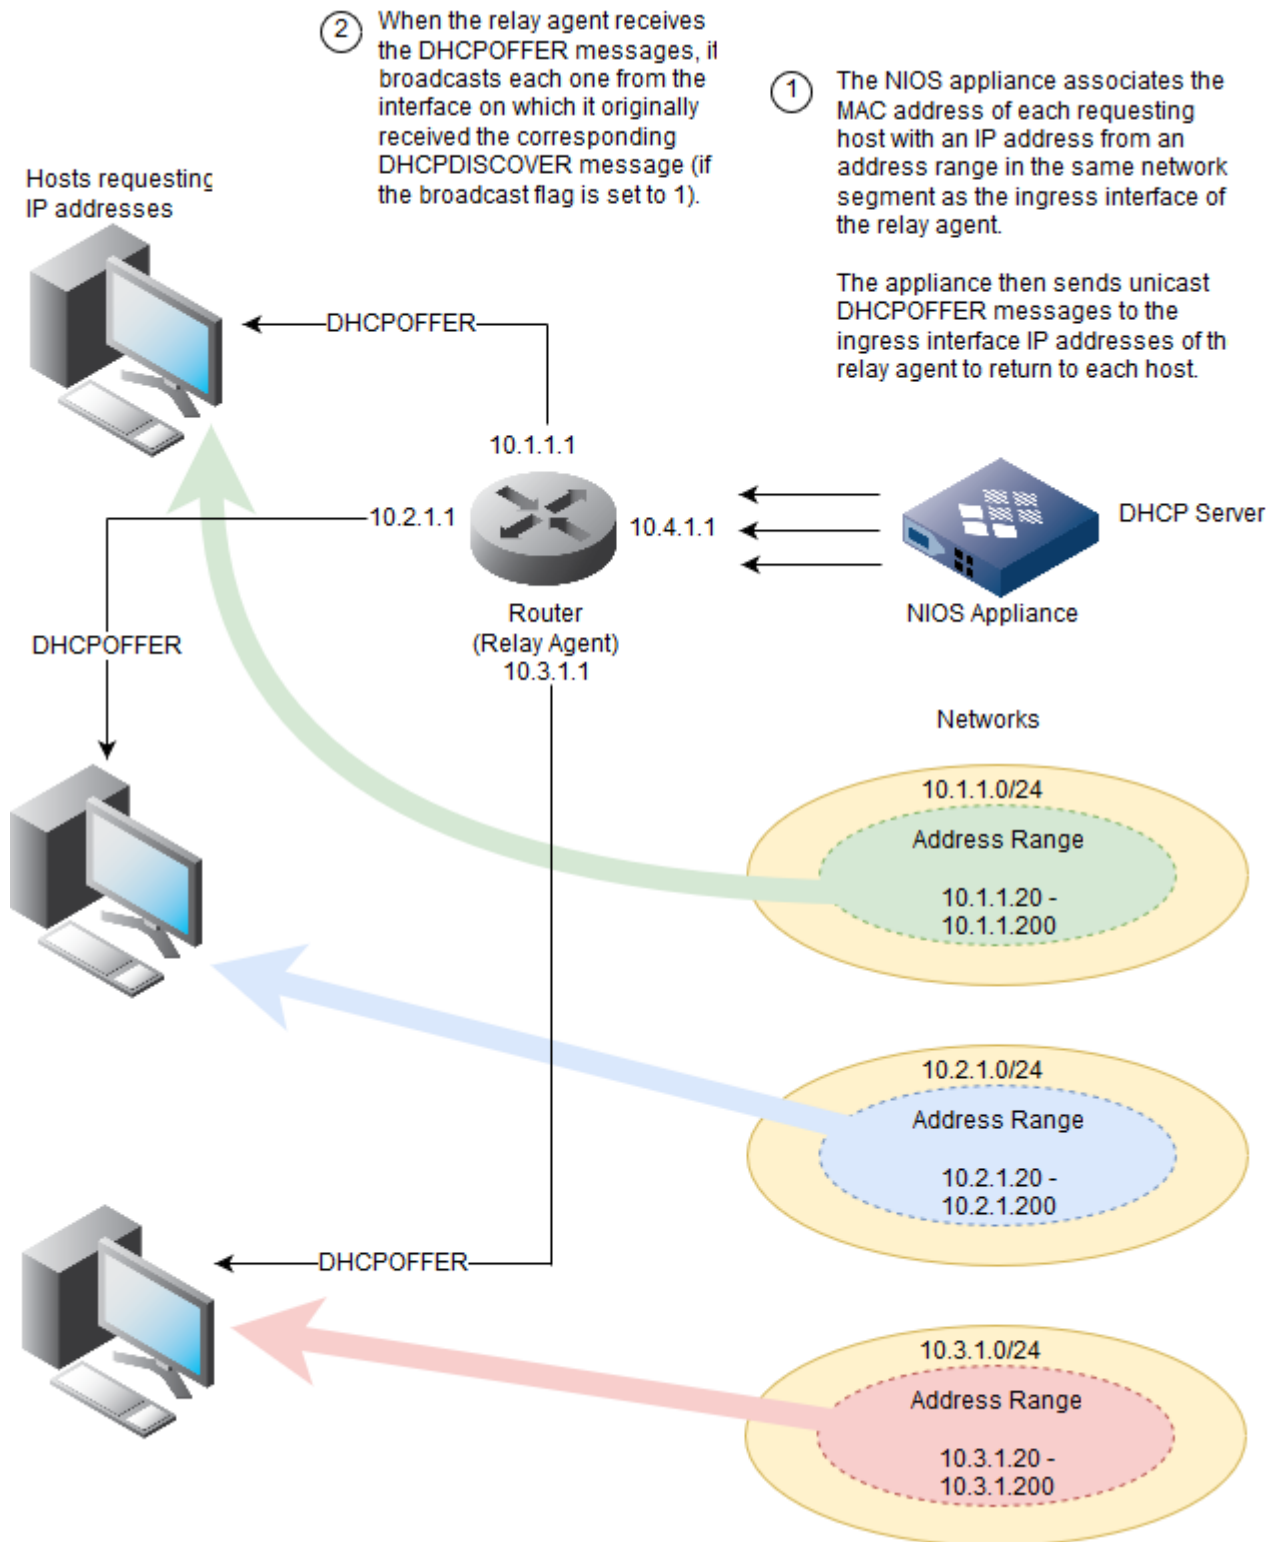
② When the relay agent receives the DHCPDISCOVER message it adds the IP address of the ingress interface to the message. It then forwards the message using IP unicast to the DHCP server.

③ When the NIOS appliance receives the DHCPDISCOVER message, it uses the ingress interface IP address on the relay agent to determine the network segment to which the host belongs. It then assigns an address to that host from the address range belonging to that subnet. Furthermore, it associates the IP address with the source MAC address of the host.



The NIOS appliance replies to DHCPREQUEST messages by sending DHCP OFFER messages through the relay agent to the requesting hosts, as shown in the following figure Requesting Addresses – DHCP OFFER Messages.

Requesting Addresses – DHCP OFFER Messages



The addressing scheme depicted in the figures Requesting Addresses – DHCPDISCOVER Messages and Requesting Addresses – DHCP OFFER Messages above, is fairly simple: each network has a single address range. Consequently,

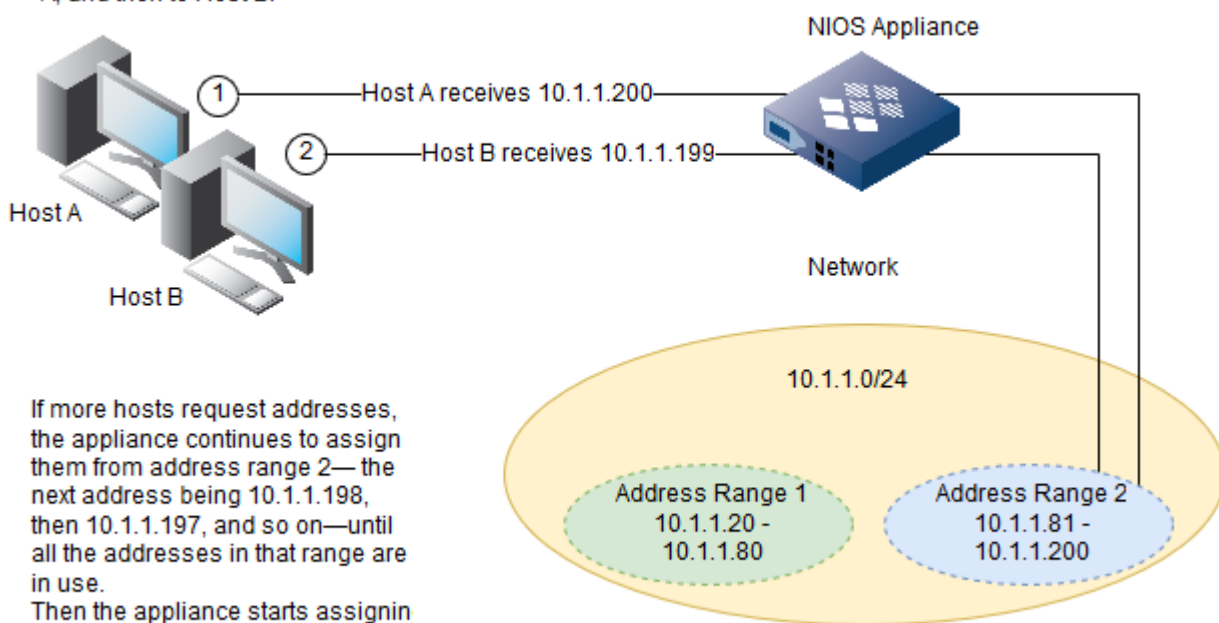
address assignments are fairly straightforward. However, if you have multiple address ranges in the same network and you want to assign addresses from specific address ranges to specific hosts, you must screen the address assignments through the use of filters.

**Note**

After the DHCP server runs for a while, it assigns leases based on when it last used addresses, and not just on their positions in the range.

### Multiple Address Ranges without Filters

The NIOS appliance assigns addresses to both hosts from the same address range—first to Host A, and then to Host B.



If more hosts request addresses, the appliance continues to assign them from address range 2—the next address being 10.1.1.198, then 10.1.1.197, and so on—until all the addresses in that range are in use. Then the appliance starts assigning addresses from address range 1, starting at 10.1.1.80, and stopping at 10.1.1.200.

## IP Address Allocation Using Filters

To control the assignment of addresses from specific address ranges to specific hosts, the NIOS appliance provides the following filters:

- A MAC address filter to which you add MAC addresses as filter criteria. For information, see [Configuring MAC Address Filters](#).
- A relay agent filter with configured circuit ID and remote ID as specified by the relay agent (DHCP option 82). For information, see [About Relay Agent Filters](#).
- An option filter in which you specify DHCP options and matching values. For information, see [Configuring Option Filters](#).
- A NAC filter in which you specify authentication results from a RADIUS authentication server group as filter criteria. For information, see [About NAC Filters](#).

When the appliance receives an address request, it checks if the request matches a filter. If it does not, the appliance assigns an address from the address range with the highest available IP address. If the request matches at least one class filter for a range, the appliance applies the following rules:



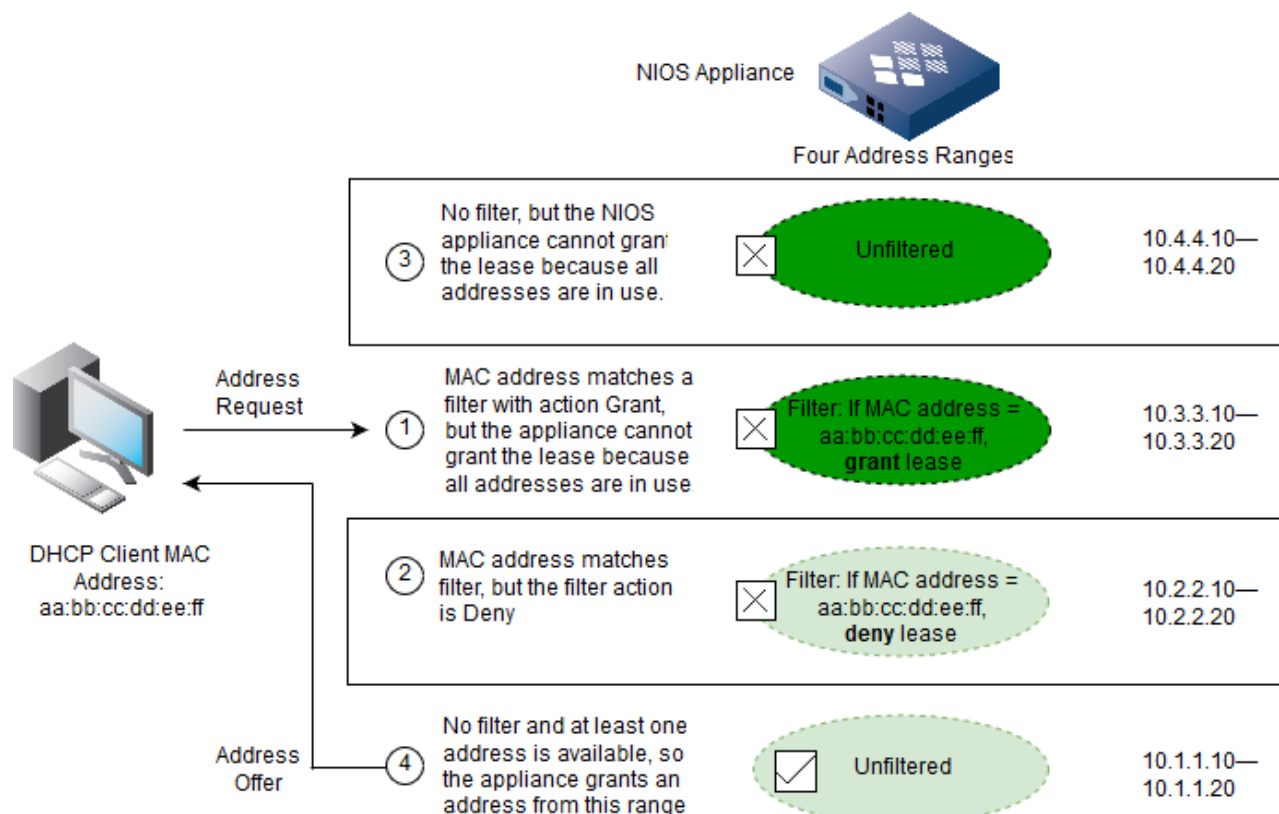
- If there are grant address filters applied to that range, the request must match one of the class filters or the appliance does not grant an address from that range.
- If there are deny address filters applied to that range, the request must not match any of the filters. If the request matches a deny filter, the appliance does not grant an address from that range.
- If an address range has a combination of grant and deny filters, the request must:
  - Match a grant filter
  - Not match a deny filter

Two rules govern the behavior of the appliance in relation to DHCP filters:

1. Depending on your filter configuration, the appliance checks if any data in an address request (such as the MAC address of the client, DHCP options 77 and 82, etc.) matches any filters applied to an address range.
2. The appliance checks for available addresses in the address ranges containing the highest addresses first. ("Highest" means closest to 255.255.255.255, and "lowest" means closest to 0.0.0.0.)

These two rules can work in coordination. For example, when the appliance receives an address request, it first checks if the request matches any filter. If it matches more than one filter assigned to different address ranges, the appliance first applies the filter that belongs to the range with the highest IP addresses. If that address does not grant an address lease (because the filter action is Deny or all address leases in that range are already in use), the appliance then applies the matching filter for the range with the next higher set of IP addresses. If the appliance still has not granted a lease from the address ranges whose filters match data in the request and there are unfiltered address ranges, the appliance attempts to assign an address from one of these ranges, again beginning with the range having the highest IP addresses. The figure DHCP Address Assignment with Multiple Filters below presents an example illustrating the sequence in which the appliance assigns addresses when a request matches a MAC address filter. For information about MAC address filters, see [Configuring MAC Address Filters](#).

Figure 31.4 DHCP Address Assignment with Multiple Filters



if	then
the appliance receives a request that matches a filter for one address range,	it applies the action specified in the filter for that address range. If it does not assign an address from that range (the action is deny or the action is grant but all addresses in that range are in use), the appliance then checks if it can assign an address from an unfiltered address range (if there are any), starting with the range with the highest addresses first, as shown in <a href="#">IP Address Allocation</a> .
the same filter applies to multiple address ranges and the appliance receives an address request matching that filter,	it checks the address range with the highest IP addresses matching that filter. If the appliance does not assign an address from that range, it checks the filtered address range with the next highest IP addresses, and so on. If it still has not assigned an address, the appliance starts checking unfiltered address ranges (if there are any), again beginning with the range with the highest address first.
multiple filters for the same address range conflict with each other (one filter grants a lease and another denies it) and a requesting client matches both filters,	the filter denying the lease takes precedence. For example, if a requesting client matches both a MAC address filter (granting a lease) and a user class filter (denying a lease) for the same address range, the appliance denies the lease. When faced with a choice to either allow or deny a lease based on equal but contradictory filters, the appliance takes the more secure stance of denying it.

## Configuring MAC Address Filters

The appliance can filter an address request by the MAC address of a requesting host. Depending on how you apply the MAC filter, the appliance can grant or deny the address request if the requesting host matches the filter criteria. You can also define DHCP options that you want to return to the matching client if the options are so configured. The client can also request specific options to be returned through DHCP option 55. The appliance returns DHCP options to matching clients based on how you apply the filters. For information, see [Applying Filters to DHCP Objects](#).

You can configure a MAC address filter or specific MAC addresses within a filter to expire after a certain amount of time has passed. Filter expiration is useful in situations where you want to keep filters running against updated MAC addresses. The permission to use the MAC addresses assigned to an IP address may become invalid after a certain period of time. For example, you can use a MAC address filter to restrict the right to use MAC addresses assigned to IP addresses for visiting guests or temporary workers. You can avoid removing invalid addresses from address filters manually by configuring the appliance to expire filters or to expire specific addresses within filters.

To apply a MAC address filter to an address range:

1. Define a MAC address filter. For information, see [Defining MAC Address Filters](#) below.
2. Add a MAC address to the filter. For information, see [Adding MAC Address Filter Items](#) below.
3. Apply the filter to a DHCP address range or range template, and specify that if the MAC address of a requesting host matches the filter definition, the appliance either grants or denies the address assignment. For information, see [Applying Filters to DHCP Objects](#).

### Defining MAC Address Filters

To define a MAC address filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab, and then expand the Toolbar and click **Add -> IPv4 MAC Address Filter**.  
or  
From any panel in the **DHCP** tab, expand the Toolbar and click **Add -> IPv4 MAC Address Filter**.
2. In the *Add IPv4 MAC Filter* wizard, complete the following:
  - **Name:** Enter a meaningful name for the filter. For example, if you want to filter address requests by department, you can name one filter "Marketing", another "Finance", and so on. The name must be unique within a specific network. If you want to specify option settings in the filter, the filter name must be unique among all MAC filters.
  - **Comment:** Enter useful information about the filter.
3. Click **Next** and complete the following to define the DHCP options to return to the matching client:
  - **Option Space:** Select an option space from the drop-down list. This field is displayed only when you have custom option spaces. The appliance uses the **DHCP** option space as the default.

- **Lease Time:** Enter the value of the lease time in the field and select the time unit from the drop-down list. The lease time applies to hosts that meet the filter criteria.

#### Options to Merge with Object Options

Click the Add icon. Grid Manager adds a new row to the table with the default **DHCP** option space and option name displayed. Complete the following:

- **Option Space:** Click the down arrow and select an option space from the drop-down list. The selected option space contains the corresponding DHCP options.
  - **Option Name:** Click the down arrow and from the drop-down list, select the DHCP option you want to return to the requesting host.
  - **Value:** Enter the value that you want the filter to return for the selected DHCP option. For example, enter the value 255.255.255.0 for the subnet-mask option.  
To add more options to the filter, click the Add icon and repeat the steps.
4. Click **Next** and complete the following to configure the expiration setting:
    - **Default MAC Address Expiration**
    - Select one of the following to configure the expiration setting for the filter:
      - **Never Expires:** Select this if you want the MAC address filter to never expire. This is selected by default.
      - **Automatically Expires in:** Select this if you want the filter to expire after a specific time frame. You can specify the time in seconds, minutes, hours, or days.  
The filter expiration time you configure here affects how long the DHCP server grants a lease to a client. It has an upper limit of 15 minutes on the lease time you configure for the Grid. For example, if both the filter expiration time and the lease time are less than 15 minutes, the appliance uses the lease time. If both the filter expiration time and lease time are greater than 15 minutes, the appliance uses the filter expiration time. If the filter expiration time is less than 15 minutes and the lease time is greater than 15 minutes, the DHCP server grants a lease for 15 minutes. If the filter expiration time is greater than 15 minutes and the lease time is less than 15 minutes, the appliance uses the lease time.
      - **Enforce Expiration Times :** Select this to enable the expiration setting.
      - **Enabled:** The filter is enabled by default. Clear the checkbox box to disable this filter.
  5. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).
  6. Save the configuration and click **Restart** if it appears at the top of the screen.

## Adding MAC Address Filter Items

To add a MAC address to a MAC address filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab, and then expand the Toolbar and click **Add -> IPv4 MAC Address Filter Item**.  
or  
From any panel in the **DHCP** tab, expand the Toolbar and click **Add -> IPv4 MAC Address Filter Item**.
2. In the *Add IPv4 MAC Address Filter Item* wizard, complete the following:
  - **MAC Address Filter:** Click **Select Filter**. In the *DHCP Filter Selector* dialog box, select the MAC address filter to which you want to add a MAC address, and then click the Select icon. If you are adding a MAC address to a filter that you have selected in the *Filters* panel, Grid Manager displays the selected filter in this field.
  - **MAC Address:** Enter the MAC address in one of the following formats: aa:bb:cc:dd:ee:ff, aa-bb-cc-dd-ee-ff, aabb.ccdd.eeff, aabbcc-ddeeff, and aabbccddeeff. The appliance displays the address in the AA:BB:CC:DD:EE:FF format. You can also enter a vendor prefix in the three hexadecimal format using the same separators supported in the MAC address format. For example, you can enter aa.bb.cc as the vendor prefix. The appliance displays AA:BB:CC.
  - **Comment:** Enter useful information about the filter item.
  - **Expiration Time:** MAC addresses in a filter stay valid until you explicitly configure them to expire. You can enable expiration for specific MAC addresses in the filter. Select one of the following:
    - **Never Expires:** Select this if you want the MAC address to never expire. This is selected by default.
    - **Expires on:** Select this and specify the **Date** and **Time** for the expiration. The fields display the current date and time. If you have already configured an expiration time for the filter, the appliance

displays the time here by adding the filter expiration time to the current time. For example, if the expiration time for the filter is two days and the current date is June 6, 2009, the appliance displays June 8, 2009 in the **Date** field.

3. Click **Next** and select one of the following to configure user registration (optional):
  - **Register as User:** Select this and enter a username in the field.
  - **Register as Guest:** Select this and enter the first name, middle name, last name, email address, and phone number of the guest user.  
The appliance displays the information you enter here in the lease viewers.
4. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).
5. Save the configuration and click **Restart** if it appears at the top of the screen.

After you define a MAC address filter and add MAC addresses to it, you can assign the filter to a DHCP range. The appliance filters IP address requests based on the filter criteria. For information, see [Applying Filters to DHCP Objects](#).

### Guidelines for Associating MAC Address Filters with DHCP Objects

The following is a list of guidelines for associating MAC Address Filters with DHCP objects:

- Associating a MAC address filter with a DHCP range on a DHCP member or failover association requires a DHCP service restart.
- Associating the same MAC address filter to a different DHCP range on the same member does not require a DHCP service restart.
- Making further changes to a MAC address filter already assigned to a range (such as changing a MAC filter policy on the range from the grant lease to the deny lease) does not require a service restart.
- A restart is required only when:
  - You add a new MAC address filter (new to the DHCP member) is added to a range on a specific DHCP member or failover association.
  - Remove a specific MAC address filter from the last range it was associated with, on a specific DHCP member or failover association.
- Changes are propagated seamlessly to the the dhcpd process memory without a service restart and the dhcpd.conf file is not updated until you restart the service.

### About Relay Agent Filters

The NIOS appliance can filter an address request by the circuit ID and remote ID of a requesting host. The filter instructs the appliance either to grant or deny an address request if the requesting host matches the filter. For information about the DHCP relay agent option, see as described in [About the DHCP Relay Agent Option \(Option 82\)](#).

Option 82 assists the agent in forwarding address assignments across the proper circuit. When a relay agent receives a DHCPDISCOVER message, it can add one or two agent IDs (circuit ID or remote ID) in the DHCP option 82 suboption fields to the message, as illustrated in the figure Relay Agent Filtering below. If the agent ID strings match those defined in a relay agent filter applied to a DHCP address range, the appliance either assigns addresses from that range or denies the request based on the configured parameters.

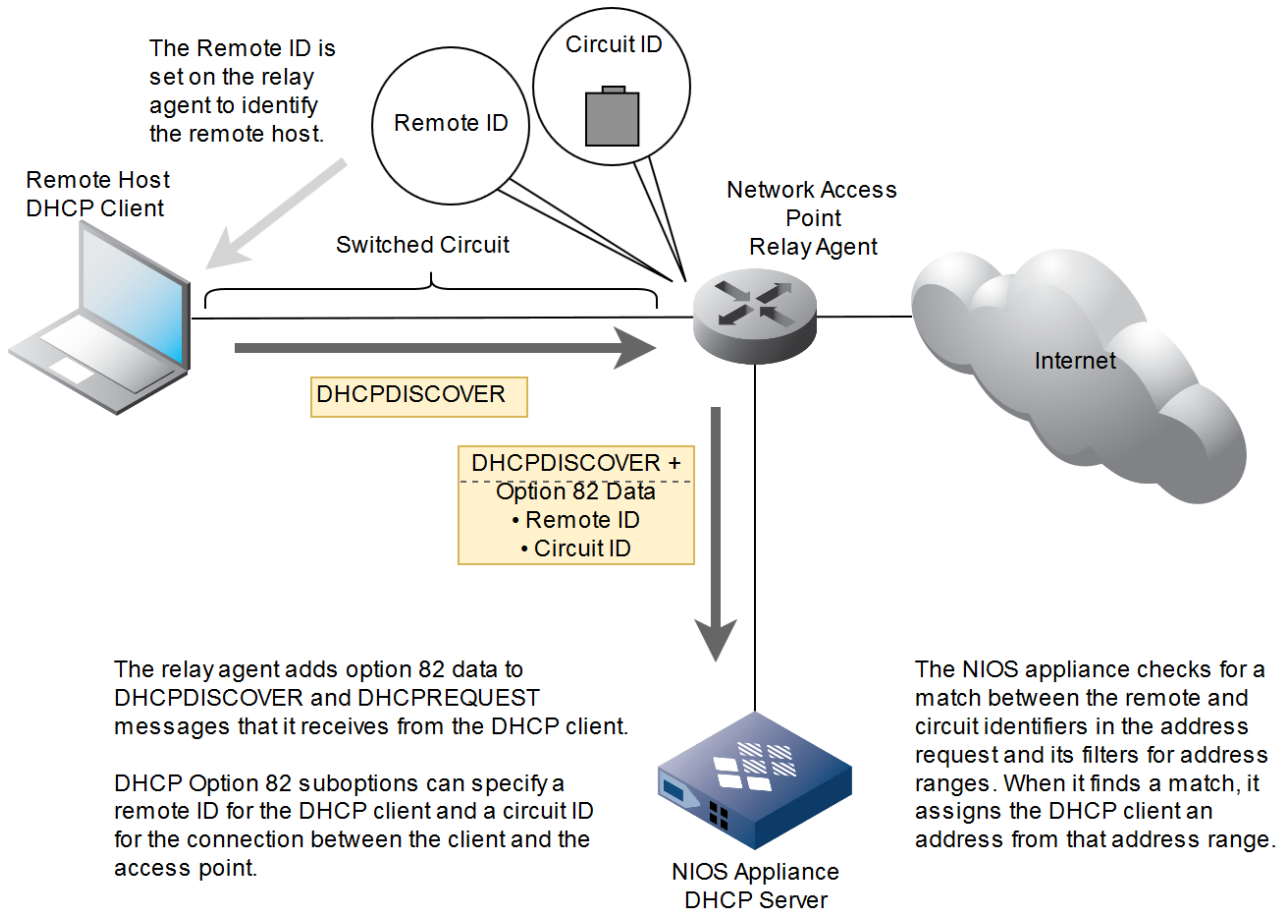
#### *Relay Agent Filtering*

Sample client ID types:

- Unit-slot-port number of network access server
- ATM virtual circuit number
- Cable data virtual circuit number

Sample remote ID types:

- User name (as prompted by network access server)
- Remote caller ATM address
- Modem ID for a cable data modem



To apply a relay agent filter to an address range:

1. Define a relay agent filter. For information, see [Defining Relay Agent Filters](#) below..
2. Apply the filter to a DHCP address range or range template, and specify that if the circuit ID or remote ID of a requesting host matches the filter definition, the appliance either grants or denies the address assignment. For information, see [Applying Filters to DHCP Objects](#).
3. Define the access privileges of limited-access admin group for relay agent filters. For information, see [Managing Administrators](#).

## Defining Relay Agent Filters

To define a relay agent filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab, and then expand the Toolbar and click **Add -> IPv4 Relay Agent Filter**.  
or  
From any panel in the **DHCP** tab, expand the Toolbar and click **Add -> IPv4 Relay Agent Filter**.
2. In the *Add IPv4 Relay Agent Filter* wizard, complete the following:
  - **Name:** Enter a meaningful name for the filter. For example, you can enter the IP address or the name of the router acting as the relay agent. The name must be unique within a specific network.

- **Comment:** Enter useful information about the filter.
3. Click **Next** to define the relay agent ID type. If you apply both ID types, the relay agent must provide both identifiers when submitting a DHCP address request.  
Select one of the following for both **Circuit ID** and **Remote ID**:
    - **Any:** Select this and the filter matches any of the circuit identifiers for remote hosts. You cannot select this for circuit ID and remote ID at the same time.
    - **Not Set:** Select this and no circuit identifier is set for remote hosts.
    - **Matches Values:** Select this and enter the circuit ID or remote ID in the field. You can enter the ID in hexadecimal format, such as 1f:cd, ab, or ef:23:56, or in string format, such as abcd or aa:gg. The appliance matches the value you enter here with the value sent by the DHCP client in counted octet sequence format. This field supports wildcard characters and regular expressions. You can also select to have an exact match or a substring match, as follows:
      - **Exact Match:** Select this to match the exact value sent by the DHCP client that contains the value you entered in the **Matches Values** field.
      - **Substring:** Select this to match a substring of the value sent by the DHCP client. Enter the **Offset** and **Length** values for the substring match, as follows:
        - **Offset:** Enter the number of characters at which the match value for the substring starts. Enter 0 to start at the beginning of the value, enter 1 for the second position, and so on. For example, when you enter 2 and have a substring value of AFTR, the appliance matches the value AFTR starting at the third character of the match value.
        - **Length:** Enter the length of the substring value. For example, if the match value is AFTR, the length is 4.
  4. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).
  5. Save the configuration and click **Restart** if it appears at the top of the screen.

After you define a relay agent filter, you can assign it to a DHCP range. The appliance responds to address requests based on the filter criteria. For information, see [Applying Filters to DHCP Objects](#).

## Configuring Option Filters

You can use option filters to classify DHCP clients and decide which DHCP options each group of clients can receive. By default, regardless of the networks in which the DHCP clients reside and whether an option filter is applied to a DHCP range or range template, all DHCP clients that match the filter criteria receive the DHCP options and values you define in the filter. You can change this configuration so the appliance does not use the filter to classify DHCP clients. For information about how to configure this, see [Defining Option Filters](#) below.

You can add DHCP options and the Hardware Operator option to an option filter. (For information about the Hardware Operator option, see [DHCP Hardware Operator](#) below.) Depending on whether the options you add to the filter are also defined at the Grid, member, network, and DHCP range levels, and whether you add the filter to the Class Filter List or Logic Filter List, the appliance either appends them to the existing options or overwrites the option values before returning them to the matching clients. For more information about how the appliance returns DHCP options, see [Adding Filters to the Logic Filter List](#).

The appliance can filter an address request by the options (such as root-server-ip-address or user-class) of the requesting host. Depending on how you apply an option filter, the appliance can grant or deny an address request if the requesting host matches the filter criteria. You can also create complex match rules that use the AND and OR logic to further define the filter criteria. When you select match rules in Grid Manager, you can preview the rules before committing them to the filter. Grid Manager provides an expression builder that automatically builds the rules after you define them.

To define an option filter and apply it to an address range:

1. Define an option filter based on either the predefined or custom DHCP options.
2. Apply the filter to a DHCP address range or range template in the Class Filter List or Logic Filter List. For information, see [Applying Filters to DHCP Objects](#).

After you define an option space and add options to it, you can set up option filters and define option values. For example, to handle two different client classes, you can define two option filters (vendor-class\_1 and vendor-class\_2) and send different option values to different clients based on the vendor-class-identifier options that you obtain from the clients.



## DHCP Hardware Operator

You can define the Hardware Operator option and add it as a match rule to an option filter. This option enables the appliance to match the hardware type and MAC address of the DHCP client, which it derives from the hardware type, hlen (hardware length) and chaddr (client hardware address) fields of the client's DHCP Discover and Renew packets. To add Hardware Operator to an option filter, fill in the fields as follows:

- In the first drop-down list, select **Hardware Operator**. Note that because it is not a DHCP Option, it does not have an actual option number.
- In the second drop-down list, select one of the following operators: **equals**, **does not equal**, **substring equals** and **substring does not equal**.  
If the operator is **substring equals** or **substring does not equal**, specify the offset and length.
- In the text field, enter the string that represents the hardware type and MAC address to match. For example, the htype value is 1 for the Ethernet hardware type. The hardware types (hrd) are defined at <http://www.iana.org/assignments/arp-parameters/arp-parameters.xml>.

This filter rule assumes that the values exist in the DHCP packets.

The following table provides examples of rules that include the Hardware Operator option. The entry in the first drop-down list for all rules is **Hardware Operator**.

### Hardware Operator Sample Rules

Rule Description	Second Drop-Down List (operator)	Text Field (string)	Offset	Length
Match a hardware type and MAC address.	equals	01:00:C0:B0:AA:BB:CC		
Match hardware type only.	substring equals	01	0	1
Match the vendor MAC prefix (first three bytes of MAC address).	substring equals	00:C0:B0	1	3

## Defining Option Filters

To define an option filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab, and then expand the Toolbar and click **Add -> IPv4/IPv6 Option Filter**.
2. or  
From any panel in the **DHCP** tab, expand the Toolbar and click **Add -> IPv4/IPv6 Option Filter**.
3. In the *Add IPv4 Option Filter* wizard, complete the following:
  - **Name:** Enter a meaningful name for the option filter. For example, you can enter Linux if you plan to use this option filter to screen Linux systems. The name must be unique within a specific network. If you want to specify option settings in the filter, the filter name must be unique among all option filters.
  - **Comment:** Enter useful information about the filter.
  - **Apply this filter as a global DHCP class:** This checkbox is selected by default. When you select this checkbox, the appliance defines a global class statement in the dhcpd configuration file for members that have DHCP enabled, regardless of whether the filter is applied to a DHCP range or range template. All DHCP clients that belong to this class receive the DHCP options and values you define in the filter. When you clear this checkbox, you cannot apply this filter to the Class Filter List of a range or range template. You cannot clear this checkbox if the filter is currently applied to a range or range template. The appliance displays an error message when you try to save this configuration.
4. Click **Next** and complete the following to add match rules:
  - In the first drop-down list, select a DHCP option.  
For example,



- If you are adding an IPv4 option filter, select **user-class(77)** for a specific user class, such as mobile users.
- If you are adding an IPv6 option filter, select **dhcp6.fqdn (39) string** for a specific domain name. The following DHCP options are not supported for IPv6 option filter:
  - dhcp6.bcms-server-a
  - dhcp6.bcms-server-d
  - dhcp6.domain-search
  - dhcp6.name-servers
  - dhcp6.nis-domain-name
  - dhcp6.nisp-domain-name
  - dhcp6.nis-servers
  - dhcp6.preference (7) 8-bit unsigned integer
  - dhcp6.rapid-commit (14) boolean
  - dhcp6.server-id (2) string
  - dhcp6.sip-servers-addresses
  - dhcp6.sip-servers-names(21) domain-list
  - dhcp6.sntp-servers
  - dhcp6.unicast(12)ip-address
- In the second drop-down list, select an operator.
 

If you select **equals** or **does not equal**, enter the value of the selected option you want the filter to match in the field.

If your operator and match value include a substring of an option value, enter the offset and length of the substring based on the following definitions:

  - **Offset:** Enter the number of characters at which the match value substring starts in the option data. Enter 0 to start at the beginning of the option data, enter 1 for the second position, and so on. For example, when you enter 2 and have a match value of MSFT, the appliance matches the value MSFT starting at the third character of the option data.
  - **Length:** Enter the length of the match value. For example, if the match value is MSFT, the length is 4.

You can do the following and repeat the filter selection steps to add another rule:

- Click **+** to add another rule at the same level.
- Click **|<** to add an **all** (logical AND) or **any** (logical OR) operator line and a parenthetical rule that is indented one level and above the first rule.
- Click **->|** to add an **all** (logical AND) or **any** (logical OR) operator line and a parenthetical rule that is indented one level.

After you add all the match rules, you can click **Preview** to view the rules that are written to the dhcpd configuration file or click **Reset** to remove the previously configured rules and start again. For information about how to use match rules, see *Using Match Rules in Option Filters* below.

5. Click **Next** and complete the following to define which DHCP options to return to the matching client:
  - **Option Space:** For an **IPv4 or IPv6 option filter** select an option space from the drop-down list. This field is not displayed if you do not have custom option spaces. The appliance uses the **DHCP** option space as the default.
  - **Lease Time:** Enter the value of the lease time in the field and select the time unit from the drop-down list. The lease time applies to hosts that meet the filter criteria.

#### Options to Merge with Object Options

Click the Add icon. Grid Manager adds a new row to the table with the default **DHCP** option space and option name displayed. Complete the following:

- **Option Space:** Click the down arrow and select an option space from the drop-down list. The selected option space contains the corresponding DHCP options that you can use as filter criteria.
  - **Option Name:** Click the down arrow and from the drop-down list, select the DHCP option you want to use as filter criteria.
  - **Value:** Enter the match value that you want the filter to use for the selected DHCP option. To add more options to the filter, click the Add icon and repeat the steps.
6. Click **Next** to define extensible attributes. For information, see *Managing Extensible Attributes*.
  7. Save the configuration and click **Restart** if it appears at the top of the screen.

## Using Match Rules in Option Filters

Each match rule you define in an option filter further defines the filter criteria of a matching client. You can add multiple match rules to an option filter. The appliance writes these rules to the dhcpd configuration file. You can also create complex match rules that use the AND and OR logic to further define the filter criteria. After you define the match rules, you can preview the rules before committing them to the filters.

For example, you can define the following rules in an option filter:

DHCP option = `vendor-class-identifier`

Substring offset = `0` (the match value starts at the beginning of the option data received from the client)

Substring length = `4` (the length of the match value MSFT)

Match value = `MSFT`

The appliance generates the following rules in the dhcpd configuration file:

```
class "microsoft-other" {  
  
    match if substring (option vendor-class-identifier,  
0, 4) = "MSFT";  
    vendor-option-space MSFT;  
}  
}
```

**DHCP option**

**Match value**

**Length of the match value**

**Substring offset**

You can also define more complex rules using the AND and OR logic as follows:

DHCP option = `vendor-class-identifier`

Match value = `infoblox2000a`

OR

DHCP option = `vendor-encapsulated-options`

Substring offset = `0` (the match value starts at the first character of the option data received from the client)

Substring length = `8` (the length of the match value infoblox)

Match value = `infoblox`

AND

DHCP option = `vendor-encapsulated-options`

Substring offset = `10` (the match value starts at the ninth character of the option data received from the client)

Substring length = 5 , the length of the match value 2000a

Match value = 2000a

The appliance generates the following rules in the dhcpd configuration file:

```
class "infoblox" {  
    match if (option vendor-class-identifier=infoblox2000a:) or  
    ((substring(option vendor-encapsulated-options,0,8)="infoblox") and  
    (substring(option vendor-encapsulated-options,10,5)="2000a")); vendor-option-  
space DHCP  
}
```

### Configuring User Class Filters

The NIOS appliance can filter DHCP address requests by user class filters. A user class indicates a category of user, application, or device of which the DHCP client is a member. User class identifiers are configured on DHCP clients and are sent during a DHCP address request operation. The client includes the user class identifier in DHCP option 77 when sending DHCPDISCOVER and DHCPREQUEST messages.

By using user class identifiers, a DHCP server can screen address requests and assign addresses from select address ranges based on the different user class identifiers it receives. For example, if you assign a user class filter named mobile to a range of addresses from 10.1.1.31–10.1.1.80, the appliance selects an address from that range if it receives an address request that includes the user class name mobile and there are still addresses available in that range. You might want mobile users to receive these addresses because you have given them shorter lease times than other, more stationary DHCP clients. See the below figure.

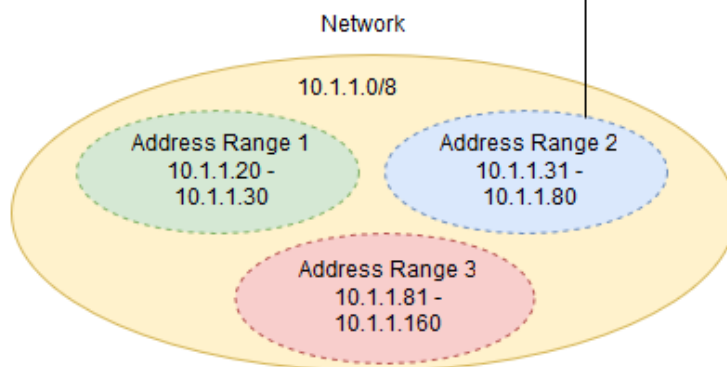
#### *Applying User Class Filtering*

The user class for laptop A is mobile. When it sends DHCPDISCOVER and DHCPREQUEST messages, it includes its user class in the DHCP option 77 field.

The NIOS appliance has a filter that screens address requests by user class. If the user class for a DHCP client is mobile, the appliance assigns it an address from address range 2.



Note: The leases for addresses in address range 2 are shorter than those for more stationary computers. The intended use for address range 2 is to provide IP addresses for mobile users who log in to the network for relatively short periods of time and, therefore, do not require longer leases.



If the NIOS appliance receives address requests with the user class mobile and there are no available addresses in address range 2 but there are available addresses in ranges 1 and 3, the appliance begins assigning addresses from address range 3 (because its addresses are higher than those in range 1). Then, if all addresses in range 3 are in use, the appliance begins assigning addresses from address range 1. If you want the appliance to assign addresses to mobile users (that is, those identified with the user class mobile) exclusively from address range 2, then you must apply user class filters for "mobile" to address ranges 1 and 3 that deny lease requests matching that user class.

### Configuration Example: Using Option Filters

The following example shows you how to create an option space, add custom options to it, create an option filter, and a match rule to filter the options so that the NIOS appliance can filter an address request by the vendor options of the requesting hosts. It can grant or deny an address request if the requesting host matches the filter.

1. Add an option space called MSFT, and then add the following options to it. For information, see [Applying DHCP Options](#).

Option name	Code	Type
root-mount-options	1	Text
root-server-ip-address	2	IP address
root-server-host-name	3	Text
root-server-path-name	4	Text

Option name	Code	Type
swap-server-ip-address	5	IP address
swap-file-path-name	6	Text
boot-file-path-name	7	Text
posix-timezone-string	8	String
boot-read-size	9	16-Bit unsigned integer

2. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab and click the Add icon.

3. In the *AddIPv4Filter* wizard, enter the filter name **i86pc**, and then select **Options** as the filter type.

4. Select **MSFT** as the option space, select an option, specify a value for it, and then add it to the **i86pc** option filter. You can select multiple options. Add the following options to the **i86pc** option filter:

Option name	Code	Type
root-server-ip-address	2	IP address
root-server-host-name	3	Text
root-server-path-name	4	Text
boot-file-path-name	7	Text

5. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab -> *filter\_name*, and then click the Add icon.

6. In the *AddIPv4MatchRule* wizard, select **i86pc** as the option filter, select **vendor-class-identifier(60)** as the matching option, and then enter **MSFT** as the matching value.

7. Add a DHCP range to the network. For information, see [Configuring IPv4 Address Ranges](#).

8. Apply the **i86pc** option filter to the DHCP address range. For information, see [Applying Filters to DHCP Objects](#).

9. Click **Restart** to restart services.

## Configuring DHCP Fingerprint Filters

The appliance can filter an address request by the DHCP fingerprint of a requesting client. Depending on how you apply DHCP fingerprint filters, the appliance can grant or deny the address request if the requesting client matches the filter criteria. Note that only superusers can add, modify, and delete DHCP fingerprint filters. Limited-access users cannot perform any DHCP fingerprint filter related tasks, though with the correct permissions they can apply DHCP fingerprint filters to DHCP ranges and range templates. For information about how to apply filters to DHCP ranges, see [Applying Filters to DHCP Objects](#).

You can define a DHCP fingerprint filter by selecting one or multiple DHCP fingerprints from the existing list of DHCP fingerprints, and then assign a grant or deny permission to the filter. You can then apply the filter to a DHCP address range, if DHCP fingerprint detection is enabled. For information about how to enable DHCP fingerprint detection, see [Enabling and Disabling DHCP Fingerprint Detection](#).

Note that once you apply a DHCP fingerprint filter to an address range, you cannot disable DHCP fingerprint detection or disable individual DHCP fingerprints that have been included in the filter. You must first delete or disable the DHCP fingerprint filter that you have applied to the address range before you can disable any fingerprint related tasks. For

information about how to delete a DHCP fingerprint filter, see [Deleting Filters](#).

On lease renewals, requesting clients must send the same DHCP fingerprint information in order for the appliance to properly grant or deny leases based on the configured DHCP fingerprint filters. For example, if a client sends option 55 in the original request but does not send the same information in the renewal request, and you have configured a DHCP fingerprint filter to grant a lease to this client, the appliance may not be able to properly grant a lease to this client.

To apply a DHCP fingerprint filter to an address range:

1. Define a DHCP fingerprint filter. For information, see [Defining DHCP Fingerprint Filters](#) below.
2. Apply the filter to a DHCP address range or range template, and specify that if the DHCP fingerprint of a requesting host matches the filter definition, the appliance either grants or denies the address assignment. For information, see [Applying Filters to DHCP Objects](#).

Associating a DHCP fingerprint filter with a DHCP range on a DHCP member or a failover association does not require a DHCP service restart. Changes are propagated seamlessly to the the dhcpd process memory without a service restart and the dhcpd.conf file is not updated until you restart the service.

## Defining DHCP Fingerprint Filters

To define a DHCP fingerprint filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab, and then expand the Toolbar and click **Add** -> **IPv4 Fingerprint Filter**.  
From any panel in the **DHCP** tab, expand the Toolbar and click **Add** -> **IPv4 Fingerprint Filter**.
2. In the *Add IPv4 Fingerprint Filter* wizard, complete the following:
  - **Name:** Enter a meaningful name for the filter. For example, if you want to filter address requests by a specific device class, you can name one filter "Gaming Console," another "Android Phones," and so on. The filter name must be unique among all DHCP fingerprint filters.
  - **Comment:** Enter useful information about the filter.
3. Click **Next** and then click the Add icon in the Select Fingerprints table. In the *Fingerprint Selector* dialog box, select the DHCP fingerprint you want to include in this filter. Click Add icon to select another DHCP fingerprint. When you select **No Match**, the appliance applies the filter to all requesting clients that do not send option 55 and option 60 or to clients that send values in option 55 and 60 that do not match any existing DHCP fingerprints.
4. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).
5. Save the configuration.

## Applying Filters to DHCP Objects

To further control how the appliance allocates IPv4 or IPv6 addresses to DHCP client requests, you can apply DHCP filters to determine the following:

- The class statements
- The address ranges from which it assigns leases
- When to grant or deny leases to the matching clients
- Which DHCP options to return to the matching clients

You can apply IPv4 or IPv6 logic filters at the Grid DHCP or Member DHCP. You can choose to keep the inherited properties or override them when you edit the IPv4 / IPv6 networks, IPv4 / IPv6 network containers, IPv4 / IPv6 network templates, IPv4 / IPv6 shared networks, IPv4 / IPv6 DHCP ranges, IPv4 / IPv6 DHCP range templates, IPv4 / IPv6 fixed addresses, IPv4 / IPv6 reservations, IPv4 / IPv6 fixed address templates, IPv4 / IPv6 reservation templates, or IPv4 / IPv6 host addresses.

## Adding Filters to the Class Filter List

You can apply any DHCP filter to the Class Filter List of a DHCP range or range template. The appliance uses the matching rules of these filters to select the address range from which it assigns a lease. You can define permissions for these filters to instruct the appliance whether to grant or deny a lease to the matching client. When you add a filter with a grant permission, the client must match the filter criteria to receive a lease. When you define a filter with a deny permission, clients that do not match the filter criteria still receive leases. Only the client that matches the filter criteria is

denied a lease.

Filters in the Class Filter List correspond to the class statement generated in the dhcpd configuration file, which is a classification of the client packet. All DHCP clients that match the option filter and relay agent filter criteria become members of the same class and are eligible to receive DHCP options for that class, regardless of the networks in which the clients reside. However, a client can only become a member of the MAC or NAC filter class when it is granted a lease from the DHCP range based on the filter criteria. Whether a client receives specific options and option values depends on the hierarchy of the options and how you apply the filters. For information about how the appliance returns DHCP options, see [Adding Filters to the Logic Filter List](#).

## Adding Filters to the Logic Filter List

The filters you add to the Logic Filter List correspond to the match rules that are written to the dhcpd configuration file. The appliance uses these filters to identify DHCP options and values to return to the matching clients. You can apply option, MAC, and NAC filters to the Logic Filter List. Note that a DHCP client is eligible to receive DHCP options defined in a filter if it matches the filter criteria. Whether the client receives specific options and their corresponding values depends on the hierarchy of the options and the list of options requested by the client through DHCP option 55. You can configure the appliance to ignore the option list requested by a matching client and return all the options that the client is eligible to receive. For information about how to ignore the option list requested by a client, see [Configuring General IPv4 DHCP Properties](#).



### Notes

- The appliance allows you to add an empty IPv4 logic filter at the end of the logic filter list, which means that you can add an IPv4 logic filter without defining DHCP options in it. In addition, you can change the order of the filters in the logic filter list.
- When a range has multiple class filters assigned to it, if any of the filters deny a lease to a client, then the client will not get a lease even if another class filter allows it.

The appliance decides which options and values to return to a client based on the following:

- If you have different DHCP options defined in a range and any DHCP filters in the Class Filter and Logic Filter lists, and these options do not overlap, the appliance merges and returns all options to the matching client. For example, a DHCP client obtains a lease from a DHCP address range (R) through an option filter in the Class Filter List (CF), which contains an option statement (O1) with a value of (S1). The appliance then matches a filter in the Logic Filter List (LF) that contains an option statement (O2) with a value of (S2). In this case, option statements O1 and O2 and their values S1 and S2 are merged and returned to the matching client.
- If there are overlapping DHCP options in a range and any DHCP filters in the Class Filter and Logic Filter lists, the values defined in the Class Filter List filters take precedence over those defined in the range and filters in the Logic Filter List. The appliance returns the option value defined in the class filters to the matching client. For example, a DHCP client obtains a lease from a DHCP address range (R) through an option filter in the Class Filter List (CF), which contains an option statement (O1) with a value of (S1). The appliance then matches a filter in the Logic Filter List (LF) that contains the same option statement (O1) with a value of (S2). In this case, the option value S1 defined in the option filter in the Class Filter List takes precedence and is returned to the DHCP client.
- When you apply option, MAC, and NAC filters to the Logic Filter List, the appliance translates their match rules into a DHCP if/elseif/else statement using the match rules of the first filter on the list as the "if" expression in the statement. Match rules in subsequent filters are translated into the "elseif" statements, and the last filter that does not contain any match rules is translated into the "else" statement. Note that a filter without any match rules can only be added as the last filter in the Logic Filter List.

For more information about how the appliance grants and denies leases to requesting clients and determines which DHCP options to return to the matching clients, see Configuration Example: Using the Class and Logic Filter Lists below. To apply filters:

1. **Grid:** From the **Data Management** tab -> **DHCP** tab, select **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox -> Edit icon.  
**Network:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network*



checkbox, and then click the Edit icon.

**DHCP Range:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr\_range* checkbox, and then click the Edit icon

**Fixed Address:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *fixed\_address* checkbox, and then click the Edit icon.

**IPv4 Reservation:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *reservation* checkbox, and then click the Edit icon.

**Host Address:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *host\_record* checkbox, and then click the Edit icon. Select the host IP address, and then click the Edit icon.

**IPv4 Network or Fixed Address Template:** From the **Data Management** tab, select the **DHCP** tab -> **Templates** tab -> (*IPv4 network or fixed address*) *template* checkbox, and then click the Edit icon.

- In the editor, click **Toggle Advanced Mode**, and then select the **Filters** tab.
- Logic Filter List:** You can keep the inherited IPv4 logic filters or override them. To override the value that has been inherited from the upper level, click **Override**. Click the Add icon to add a filter to match a client based on the match rules defined in the filter.

If you have only one configured DHCP filter, the appliance displays the filter in the table. Otherwise, in the *DHCP Filter Selector* dialog box, click the filter you want to add. Use SHIFT+click and CTRL+click to select multiple filters.

- Complete the following to add the **Class Filter** to a DHCP address range:
  - Click the Add icon to add a filter to identify the class of a matching client, and to grant or deny a lease to a client. For more information, see Adding Filters to the Class Filter List above.

If you have only one configured DHCP filter, the appliance displays the filter in the table. Otherwise, in the *DHCP Filter Selector* dialog box, click the filter you want to add. Use SHIFT+click and CTRL+click to select multiple filters.

For each filter you add, click the **Action** column and select one of the following from the drop-down list:

- Grant lease:**
    - For MAC address filters: Select this to assign an IP address from the address range to a requesting host whose MAC address matches the MAC address in the filter.
    - For relay agent filters: Select this to assign an IP address from the address range when one or both of the relay agent identifiers of the requesting host match the filter criteria.
    - For option filters: Select this to assign an IP address from the address range to a requesting host whose DHCP options match the DHCP options and match rules defined in the filter.
    - For NAC filters: Select this to assign an IP address from the address range to a requesting host based on the authentication results from a RADIUS authentication server group.
    - For DHCP fingerprint filters: Select this to grant a lease from the address range to a requesting host based whose DHCP fingerprint matches the DHCP fingerprint in the filter.
  - Deny lease:**
    - For MAC address filters: Select this to deny an address request from a host whose MAC address matches a MAC address in the filter.
    - For relay agent filters: Select this to deny an address request when one or both relay agent identifiers match the filter criteria in the filter.
    - For option filters: Select this to deny an address request from a host whose DHCP options match the options and match rules in the filter.
    - For NAC filters: Select this to deny an address request from a host based on the authentication results from a RADIUS authentication server group.
    - For DHCP fingerprint filters: Select this to deny a lease request when the DHCP fingerprint of the requesting host matches the DHCP fingerprint in the filter.
- The appliance uses filters in both the Class Filter and Logic Filter lists to determine the DHCP options and values it returns to the matching clients.



#### Note

You can only add a filter that does not contain any match rules as the last filter in the Logic Filter List.

- Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuration Example: Using the Class and Logic Filter Lists

The following example shows you how to define DHCP filters and apply them to the class and logic filter lists. It also shows you the DHCP configuration file that is generated based on the configuration.

In this example, you first define a MAC filter, two option filters (one without match rules), and a NAC filter, and then apply the MAC filter to the Class Filter List and the other filters to the Logic Filter List of the address range 10.34.34.6 - 10.34.34.55.

1. Configure and save a MAC filter as follows. For more information, see [Defining MAC Address Filters](#).
  - a. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab, and then expand the Toolbar and click **Add** -> **IPv4 MAC Address Filter**.
  - b. In the *Add IPv4 MAC Filter* wizard, complete the following:
    - **Name**: Enter **MAC1**.
  - c. Click **Next** and complete the following to define the DHCP options to return to the matching client:
    - **Lease Time**: Enter **1234** and select **seconds** from the drop-down list.  
**Options to Merge with Object Options**: Click the Add icon. Grid Manager adds a new row to the table with the default DHCP option space and option name displayed. Complete the following:
      - **Option Name**: Click the down arrow and select **log-server(7)** from the drop-down list.
      - **Value**: Enter **10.34.34.3** as the value for the log-server option that is sent to the client in the OFFER/ACK message.
  - d. Save the configuration.
2. Add a MAC address filter item as follows. For more information, see [Adding MAC Address Filter Items](#).
  - a. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab, and then expand the Toolbar and click **Add** -> **IPv4 MAC Address Filter Item**.
  - b. In the *Add IPv4 MAC Address Filter Item* wizard, complete the following:
    - **MAC Address Filter**: Click **Select Filter**. In the *DHCP Filter Selector* dialog box, click **MAC1**.
    - **MAC Address**: Enter **AB:DE:CC:DD:EE:01** as the MAC address.
  - c. Save the configuration.
3. Configure and save an option filter with match rules as follows. For more information, see [Defining Option Filters](#).
  - a. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab, and then expand the Toolbar and click **Add** -> **IPv4 / IPv6 Option Filter**.
  - b. In the *Add IPv4 Option Filter* wizard, complete the following:
    - **Name**: Enter **Option1**.
  - c. Click **Next** and complete the following to add match rules:
    - In the first drop-down list, select **vendor-class-identifier**.
    - In the second drop-down list, select **substring equals**, and then enter the following:
      - **Offset**: Enter **0** to match the value starting at the first character of the option data.
      - **Length**: Enter **4**.
      - Enter **MSFT** as the matching value.
    - Click **Preview** and the appliance displays the expression: **(vendor-class-identifier,0,4="MSFT")**.
  - d. Click **Next** and complete the following to define the DHCP options to return to the matching client:  
**Options to Merge with Object Options**: Click the Add icon. Grid Manager adds a new row to the table with the default DHCP option space and option name displayed. Complete the following:
    - **Option Name**: Click the down arrow and from the drop-down list, select **time-server(4)**.
    - **Value**: Enter **10.34.34.2** as the value for the time-server option that is sent to the client in the OFFER/ACK message.
  - e. Save the configuration.
4. Configure and save another option filter without match rules as follows:
  - a. In the *Add IPv4 Option Filter* wizard, complete the following:

- **Name:** Enter **Option2**.
- b. Click **Next**. Do not define any match rules.
  - c. Click **Next** again and complete the following to define the DHCP options to return to the matching client:
 

**Options to Merge with Object Options:** Click the Add icon. Grid Manager adds a new row to the table with the default DHCP option space and option name displayed. Complete the following:

    - **Option Name:** Click the down arrow and from the drop-down list, select `domain-name(6)`.
    - **Value:** Enter `www.infoblox.com`.
  - d. Save the configuration.
5. Configure and save a NAC filter as follows. For more information, see [Defining NAC Filters](#).
    - a. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab, and then expand the Toolbar and click **Add** -> **IPv4 NAC Filter**.
    - b. In the *AddFilter* Wizard, complete the following and click **Next**:
      - **Name:** Enter `NAC1`.
    - c. Create a rule as follows:
      - In the first drop-down list, select **Compliance State**.
      - In the second drop-down list, select **equals**.
      - In the third drop-down list, select **Compliant**.

Click **Preview** and the appliance displays the expression:  
`( Sophos.ComplianceState="Compliant" )`.
    - d. Click **Next** and complete the following to define DHCP options:
      - **Lease Time:** Enter **1000** and select `seconds` from the drop-down list.

**Options to Merge with Object Options:** Click the Add icon. Grid Manager adds a new row to the table with the default DHCP option space and option name displayed. Complete the following:

      - **Option Name:** Click the down arrow and from the drop-down list, select `cookies-servers(8)`.
      - **Value:** Enter `10.34.34.5`.
    - e. Save the configuration.
  6. Apply the filters to the address range as follows.
    - a. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> `10.34.34.6-10.34.34.55` checkbox, and then click the Edit icon.
    - b. In the *DHCP Range* editor, click **Toggle Advanced Mode**.
    - c. Click the **Filters** tab and complete the following:
 

**Class Filter List:** Click the Add icon and add **MAC1** as a class filter. Click the **Action** column and select **Grant lease** from the drop-down list.

**Logic Filter List:** Click the Add icon and add **Option1**, **NAC1**, and **Option2** respectively as logic filters
    - d. Save the configuration.
 

The appliance generates the following information in the DHCP configuration file based on the filter configuration in this example:

```
# MAC filter "MAC1"

class "MAC1" {
    default-lease-time 1234;
    min-lease-time 1234;
    max-lease-time 1234;
    option log-servers 10.34.34.3;
```

```

}
# NAC filter "NAC1"
    {option sophos.compliance
    state="compliant"
}
subnet 10.34.34.0 netmask 255.255.255.0 {
    pool {
        infoblox-range 10.34.34.6 10.34.34.55;
        range 10.34.34.6 10.34.34.55;
        option routers 10.34.34.1;
        # INFOBLOXMACFILTERDEBUGINFO: allow members of "MAC1";
        if (substring(option vendor-class-identifier,0,4)="MSFT") {
            # Option filter "Option1"
            option time-servers 10.34.34.2;
        }
    }
elseif (option Sophos.ComplianceState="Compliant") {
    # NAC filter "NAC1"
    default-lease-time 1000;
    min-lease-time 1000;
    max-lease-time 1000;
    option cookie-servers 10.34.34.5;
}
else {
    # Option filter "Option2"
    default-lease-time 2500;
    min-lease-time 2500;
    max-lease-time 2500;
    option domain-name "www.infoblox.com"; }

```

}

Depending on client requests and the matching criteria, the following scenarios can happen in this example:

If the requesting client matches the MAC1 and Option1 filters, the appliance returns the following:

- Lease time = 1234 seconds (from the MAC filter)
- Returned options:
  - Router(3) with a value of 10.34.34.1 (from the address range)
  - Log-server(7) with a value of 10.34.34.3 (from the MAC filter MAC1)
  - Time-server(4) with a value of 10.34.34.2 (from the option filter Option1)

If the requesting client matches the MAC1 and NAC1 filters, the appliance returns the following:

- Lease time = 1234 seconds (from the MAC filter MAC1)
- Returned options:
  - Router(3) with a value of 10.34.34.1 (from the address range)
  - Log-server(7) with a value of 10.34.34.3 (from the MAC filter MAC1)
  - Cookie-server(8) with a value of 10.34.34.5 (from the NAC filter NAC1)

If the client matches the MAC1 filter, but not the Option1 or NAC1 filters, the appliance returns the following:

- Lease time = 1234 seconds (from the MAC filter)
- Returned options:
  - Router(3) with a value of 10.34.34.1 (from the address range)
  - Log-server(7) with a value of 10.34.34.3 (from the MAC filter MAC1)
  - Domain-name(6) with a value of www.infoblox.com (from the option filter Option2)

If the requesting client does not match the MAC1 filter, no lease is granted.

## Managing DHCP Filters

You can do the following to manage DHCP filters:

- [Modifying DHCP Filters](#)
  - [Modifying MAC Address Filter Items](#)
- [Viewing DHCP Filters](#)
  - [Viewing MAC Address Filter Items](#)
- [Deleting Filters](#)

### Modifying DHCP Filters

To modify a filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab -> *filter\_name* checkbox, and then click the Edit icon.
2. For a MAC address filter, the *DHCP MAC Filter* editor provides the following tabs from which you can modify information:
  - **General**: Modify the fields as described in [Defining MAC Address Filters](#).
  - **DHCP Options**: Add or delete DHCP options. For information, see [Defining MAC Address Filters](#).
  - **Extensible Attributes**: Add or delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).
  - **Permissions**: This tab appears only if you belong to a superuser admin group. For information, see [About Administrative Permissions](#).
- For a relay agent filter, the Relay Agent Filter editor provides the following tabs from which you can modify information:
  - **General**: Modify the fields as described in [Defining Relay Agent Filters](#).

- **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).
- For an option filter, the Option Filter editor contains the following tabs from which you can modify information:
  - **General:** Modify the fields as described in [Defining Option Filters](#).
  - **Rules:** Modify the match rules as described in [Defining Option Filters](#).
  - **DHCP Options:** Modify option spaces and DHCP options in the **Basic** tab as described in [Defining Option Filters](#). You must define the **PXE Lease Time** in the **Advanced** tab.
  - **BOOTP:** Modify BOOTP settings as described in [Configuring IPv4 BOOTP and PXE Properties](#).
  - **Extensible Attributes:** Add or delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).
- For a DHCP fingerprint filter, the Add IPv4 Fingerprint Filter editor provides the following tabs from which you can modify information:
  - **General:** Modify general information, such as the name and device class, as described in [Defining DHCP Fingerprint Filters](#).
  - **Fingerprints:** Add or delete DHCP fingerprints as described in [Defining DHCP Fingerprint Filters](#).
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with the DHCP fingerprint filter. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).
- For a NAC filter, the NAC Filter editor contains the following tabs from which you can modifying information:
  - **General:** Modify the name and comment.
  - **Rules:** Modify the rules as described in [Defining NAC Filters](#).
  - **DHCP Options:** Add or delete DHCP options. For information, see [Defining NAC Filters](#).
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).

3. Save the configuration and click **Restart** if it appears at the top of the screen.

You can modify the MAC address filter items and match rules for corresponding MAC address filters and option filters. For information, see [Modifying MAC Address Filter Items](#) and [Viewing DHCP Filters](#) below.

### Modifying MAC Address Filter Items

To modify a MAC address filter item:

1. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab -> *filter\_name* -> *mac\_filter* checkbox, and then click the Edit icon.
2. The *MAC Address Filter Item* editor contains the following tabs from which you can edit data:
  - **General:** Modify the fields as described in [Adding MAC Address Filter Items](#).
  - **Registration:** Modify registration settings as described in [Adding MAC Address Filter Items](#).
  - **Extensible Attributes:** Add or delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).
3. Save the configuration and click **Restart** if it appears at the top of the screen.

### Viewing DHCP Filters

To view DHCP filters:

1. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab.
2. Grid Manager displays the following for each filter:
  - **Name:** The name of the filter.
  - **Filter Type:** The filter type.
  - **Comment:** The information about the filter.

- **Site:** The location to which the filter belongs. This is one of the predefined extensible attributes.

## Viewing MAC Address Filter Items

To view a list of MAC addresses in a specific MAC address filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab -> *filter\_name*.
2. Grid Manager displays the following:
  - **MAC Address:** The DHCP fingerprint information of client's device. This field is displayed only when users use captive portal for authentication. MAC address assigned to the filter.
  - **Username:** Grid Manager displays the username to which the MAC address belongs in the lease viewers.
  - **Comment:** The information you entered about the filter item.
  - **Expiration Time:** The expiration time you configured for the MAC address.
  - **Fingerprint:** The DHCP fingerprint information of client's device. This field is updated when users use Captive Portal for authentication.
  - **Site:** The location to which the filter belongs. This is one of the predefined extensible attributes.

## Deleting Filters

You can delete a filter that is not currently assigned to a DHCP range. You can also remove a filter from a DHCP range, and then delete the filter if it is not in use.

To delete a filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab -> *filter\_name*, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes**.

The appliance puts the deleted filters in the Recycle Bin, if enabled. You can later restore the filter if needed.

To schedule this task, click the Delete icon -> **Schedule Delete**. In the *Schedule Deletion* dialog box, click **Delete Later**, and then specify a date, time, and time zone.

## Authenticated DHCP

This feature provides the ability to control access to your IPv4 networks. (This feature does not support IPv6 networks.) You can divide a network into segments for unauthenticated, authenticated and guest users, and the DHCP server assigns clients to the appropriate segment based on their MAC addresses and authentication credentials.

For example, you can divide a network into one or more production segments for valid employees and systems, a guest segment with access only to the Internet and/or limited public servers, and a quarantine segment with access to a captive portal only. A captive portal is a web page that can provide an option to register as an authenticated user or as a guest.

On a member DHCP server, configure DHCP ranges for each access level — quarantine, authenticated, and guest — and create MAC address filters for the DHCP ranges. You can use DHCP options and Access Control Lists (ACLs) on your routers and firewall policies to define the appropriate services for each access level. On another Grid member, configure the captive portal and specify the authentication server group that authenticates the users. You can configure an authentication server group for external servers running RADIUS, LDAP, or Active Directory (AD).

When a DHCP client first sends a request for an IP address, the DHCP server offers an IP address from the quarantine range and directs the client to the captive portal, where the user can register either as an authenticated user or as a guest. When users sign in as guests or are successfully authenticated, the member automatically adds their MAC addresses to the appropriate MAC address filters and assigns addresses out of the appropriate address range.

This section includes the following topics:

- [DHCP Authentication Process](#)
- [Configuring DHCP Authentication](#)
- [Configuring Authentication Server Groups](#)
- [Configuring Captive Portal](#)
- [Defining IPv4 Network and DHCP Ranges](#)



- [Defining MAC Address Filters](#)
- [Using the Captive Portal Wizard](#)
- [Adding and Modifying the Filters and Associations](#)
- [Monitoring DHCP Authentication](#)
- [Configuration Example: Configuring Authenticated DHCP](#)
- [NAC Integration](#)
- [Configuring NAC with RADIUS Servers](#)
- [About Authentication Servers](#)
- [Configuring DHCP Ranges](#)
- [About NAC Filters](#)

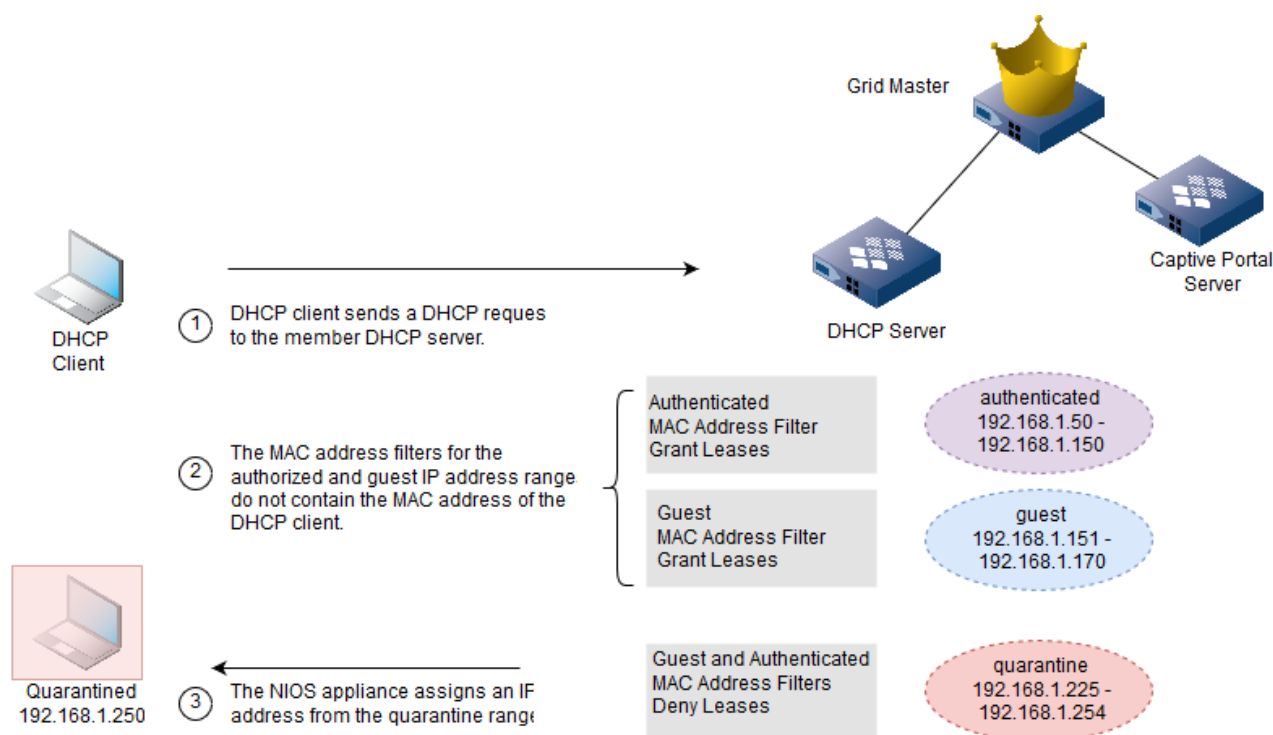
## Related topic

[DHCP Authentication Process](#)

## DHCP Authentication Process

This section illustrates the DHCP authentication process. As illustrated in the following figure [Quarantining an Unauthenticated DHCP Client](#), the DHCP authentication process begins when a DHCP client attempts to connect to the network. The member DHCP server checks if the MAC address of the DHCP client matches a MAC address in the guest or authenticated MAC address filters. If the member does not find a match, it assigns an IP address from the quarantine range to the DHCP client. When the client tries to access a web site, it is redirected to the captive portal page.

*Stage 1: Quarantining an Unauthenticated DHCP Client*



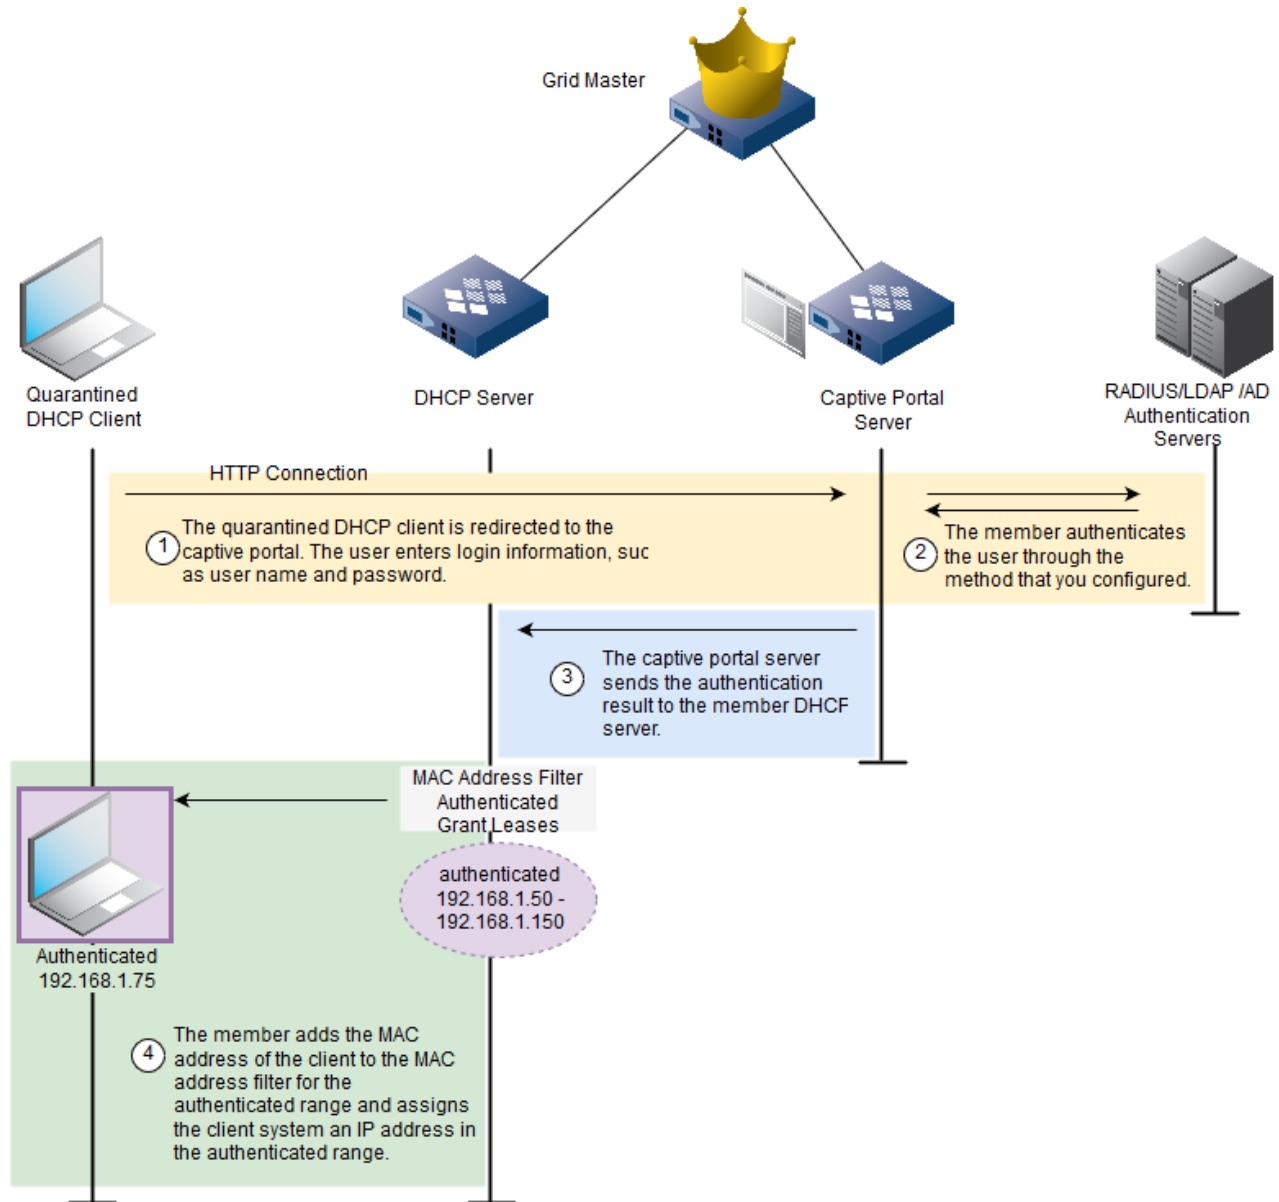
Note that the quarantine range in the figure [Quarantining an Unauthenticated DHCP Client](#) above contains MAC address filters to deny leases in the quarantine range to DHCP clients with MAC addresses that match those in the Guest and Authenticated MAC address filters.

When the client connects to the captive portal IP address through its web browser, the user can register and continue the

authentication process to obtain an IP address from the authenticated DHCP range, or register as a guest and obtain an IP address from the guest DHCP range.

If the user chooses to continue the authentication process, as shown in *Authenticating the User* below, the member authenticates the user with the authentication service that you configured, which can be RADIUS, LDAP, or AD.

*Stage 2a: Authenticating the User*

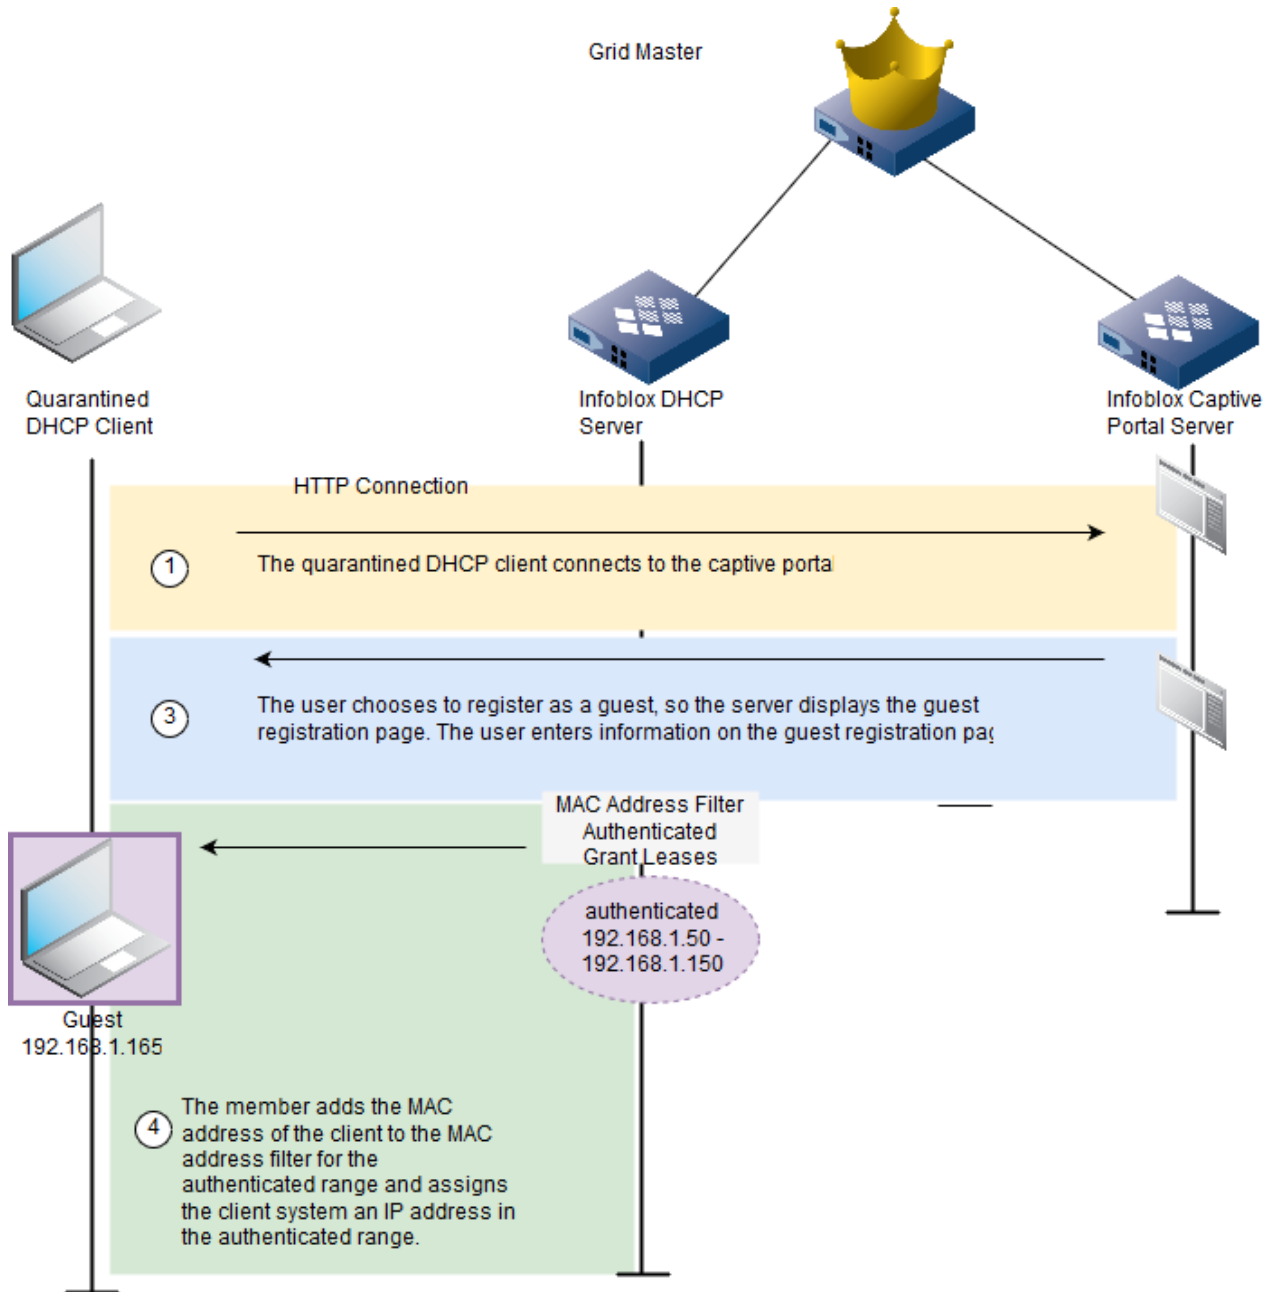


After the client successfully passes the authentication stage, the appliance stores the MAC address of the client in the MAC address filter for the authenticated range. When the client tries to renew its IP address, it receives a new IP address from the authenticated DHCP range.

Note that if the MAC address filter has an expiration period, the member automatically deletes expired MAC addresses from the filter. Therefore, if a DHCP client tries to renew its IP address after the expiration period, the client is redirected to the captive portal because its MAC address is no longer in the MAC address filter. For more information, see [Defining MAC Address Filters](#).

If the user chooses to sign in as a guest, as shown in Registering as a Guest below, the user can fill in the guest registration page provided by the captive portal.

*Stage 2b: Registering as a Guest*



After the user signs in as a guest, the appliance stores the MAC address of the client in the MAC address filter for the guest range. When the DHCP client tries to renew its IP address, it receives a new IP address from the guest DHCP range, unless the MAC address of the client expired and was removed from the filter. In this case, the DHCP client is redirected to the captive portal.

Related topic

[Authenticated DHCP](#)

## Configuring DHCP Authentication

Following are the tasks to configure the DHCP Authentication feature:

1. Configure the authentication server group which the captive portal uses to authenticate DHCP clients. For more information, see [About Authentication Server Groups](#).  
If the captive portal is used to register guest users and does not authenticate users, then you do not have to configure an authentication server group.
2. Configure the captive portal properties and associate the captive portal with the authentication server group. For more information, see [Configuring Captive Portal Properties](#).
3. Optionally, customize the captive portal interface and guest registration page, as described in [Customizing the Captive Portal Interface](#). Additionally, if you enabled SSL encryption, upload the required certificates, as described in [Managing Captive Portal Certificates](#).
4. Enable the captive portal, as described in [Starting the Captive Portal Service](#).
5. Configure the network and a DHCP range for quarantine DHCP clients. Configure DHCP ranges for authenticated and guest DHCP clients, depending on whether you are allowing either one or both types of users to access your network. For information about configuring these DHCP ranges, see [Defining the IPv4 Network and DHCP Ranges](#).
6. Run the *Captive Portal* wizard to create MAC address filters for the quarantine range and for the authenticated, and guest DHCP ranges, if configured; and to associate the captive portal server with the member that serves the DHCP ranges. To accomplish these tasks and set other properties, see [Using the Captive Portal Wizard](#). Alternatively, you can perform these tasks separately or modify the configured properties, as described in [Adding and Modifying the Filters and Associations](#).
7. Enable the DHCP service. For more information, see [Starting DHCP Services on a Member](#).

For information about monitoring the captive portal and the DHCP service, see [Monitoring DHCP Authentication](#).

## Configuring Authentication Server Groups

Create an authentication server group if you want the captive portal server to authenticate users when they register. You can create an authentication server group with RADIUS servers, LDAP servers, or Active Directory servers, and then associate the group with the member that runs the captive portal and sends the authentication requests. You can associate an authentication server group with multiple captive portals, but you can associate a captive portal with only one authentication server group.

The following sections provide instructions for creating a RADIUS authentication server group, an AD authentication server group and an LDAP server group:

- [Configuring a RADIUS Authentication Server Group](#)
- [Configuring an Active Directory Authentication Server Group](#)
  - [Managing Multiple Domain Controllers](#)
- [Configuring an LDAP Server Group](#)

### Configuring a RADIUS Authentication Server Group

You can add multiple RADIUS servers to an authentication server group and prioritize them. When the member sends an authentication request, it always selects the first RADIUS server in the list. It only sends authentication requests to the next server on the list if the first server goes down.

To configure the RADIUS authentication server group to which a captive portal server sends authentication requests:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.

2. Expand the Toolbar and click **Add** -> **RADIUS Service**.
3. In the *Add RADIUS Authentication Service* wizard, complete the following:
  - **Name:** Enter the name of the server group.
  - **RADIUS Servers:** Click the Add icon and enter the following:
    - **Server Name or IP Address:** Enter the RADIUS server FQDN or IP address.
    - **Comment:** You can enter additional information about the server.
    - **Authentication Port:** The destination port on the RADIUS server. The default is 1812.
    - **Authentication Type:** Select the authentication method of the RADIUS server from the drop-down list. You can specify either PAP (Password Authentication Protocol) or CHAP (Challenge Handshake Authentication Protocol). The default is PAP.
    - **Shared Secret:** Enter the shared secret that the member DHCP server and the RADIUS server use to encrypt and decrypt their messages. This shared secret must match the one you entered on the RADIUS server and must be between 4 and 64 characters (inclusive) in length.
    - **Connect through Management Interface:** Select this to enable the member to use its MGMT port to communicate with just this server.
    - **Disable server:** Select this to disable the RADIUS server if, for example, the connection to the server is down and you want to stop the DHCP server from trying to connect to this server.
    - Click **Test** to validate the configuration and check that the Grid Master can connect to the RADIUS server. Before you can test the configuration though, you must specify the authentication and accounting timeout and retry values.  
If the Grid Master connects to the RADIUS server using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the RADIUS server, the appliance displays a message indicating an error in the configuration.
    - Click **Add** to add the RADIUS server to the group.

When you add multiple RADIUS servers to the list, you can use the up and down arrows to change the position of the servers on the list. The member DHCP server connects to the RADIUS servers in the order they are listed.

  - **Authentication Timeout:** The time that the member DHCP server waits for a response from a RADIUS server before considering it unreachable. You can enter the time in milliseconds or seconds. The maximum is 10 seconds.
  - **Retries:** The number of times the member DHCP server retries connecting to a RADIUS server before it considers the server unreachable. The default is five.
  - **Accounting Timeout:** The time that the member DHCP server waits for a response from a RADIUS server before considering it unreachable. You can enter the time in milliseconds or seconds. The maximum is 10 seconds.
  - **Retries:** The number of times the member DHCP server retries connecting to a RADIUS server before it considers the server unreachable. The default is five.
  - **Recovery Interval:** Specifies the duration of time a RADIUS server stays inactive after being down, before becoming eligible to have RADIUS requests sent to it. The recovery interval starts when a RADIUS server is first discovered to be down.
  - **Comment:** You can enter additional information about the server group.
  - **Disable:** Select this to disable the authentication server group.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

### Configuring an Active Directory Authentication Server Group

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Click the **Active Directory Services** subtab and click the Add icon.
3. In the *Add Active Directory Authentication Service* wizard, complete the following:
  - **Name:** Enter a name for the service.
  - **Active Directory Domain:** Enter the AD domain name.
  - **Domain Controllers:** Click the Add icon and complete the following to add an AD domain controller:
    - **Server Name or IP Address:** Enter the FQDN or the IP address of the AD server that is used for authentication.
    - **Comment:** Enter additional information about the AD server.
    - **Authentication Port:** Enter the port number on the domain controller to which the member sends authentication requests. The default is 389.

- **Encryption:** Select **SSL** from the drop-down list to transmit through an SSL (Secure Sockets Layer) tunnel. When you select SSL, the appliance automatically updates the authentication port to 636. Infoblox strongly recommends that you select this option to ensure the security of all communications between the member and the AD server. If you select this option, you must upload a CA certificate from the AD server. Click **CA Certificates** to upload the certificate. In the *CA Certificates* dialog box, click the Add icon, and then navigate to the certificate to upload it.
  - **Connect through Management Interface:** Select this so that the member uses the MGMT port for administrator authentication communications with just this AD server.
  - **Disable server:** Select this to disable an AD server if, for example, the connection to the server is down and you want to stop the Grid member from trying to connect to this server.
  - Click **Test** to test the configuration. If the Grid member connects to the domain controller using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the server, the appliance displays a message indicating an error in the configuration.
  - Click **Add** to add the domain controller to the group.
  - **Timeout(s):** The number of seconds that the Grid member waits for a response from the specified authentication server. The default is 5.
  - **Comment:** Enter additional information about the service.
  - **Disable:** Select this to retain an inactive AD authentication service profile.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Managing Multiple Domain Controllers

You can create multiple domain controllers on the Microsoft server and associate it with a single Active Directory Domain, which is synchronized by multiple Grid members within the same Grid:

- Synchronization happens from only one Grid member that is a Microsoft server, which is referred as the master, at a time in a given network view.
- If the master Microsoft server fails the synchronization three times in a row, it loses its master status.
- The Grid Master will periodically ensure that for each Active Director Domain there is one Microsoft server with the master status. Otherwise, the appliance selects a new master, based on the following:
  - Microsoft server that has been designated as the Grid Master for the longest time previously.
  - Microsoft server configured in read/write synchronization mode.
- The master server is indicated within the Active Directory Domain object.
- Synchronization mode changes based on the master.
- When synchronization happens in two separate network views, respective Grid members synchronize data simultaneously.

## Configuring Captive Portal

The captive portal can be used to register users for authentication, guest users, or both types of users. When a DHCP client attempts to connect to the network and its MAC address is not in any of the configured MAC filters, the member DHCP server assigns it an IP address in the quarantine range. When the quarantined client tries to reach any web site, it is redirected to the captive portal. The captive portal runs a limited DNS server that is used solely to redirect queries to the captive portal web interface.

You can enable the captive portal as a service on any Grid member, except the Grid Master or Grid Master candidate. The Grid member that runs the captive portal cannot run any other service, such as DHCP and DNS. Note that the limited DNS service that the captive portal runs is different from the full-scale DNS service on an Infoblox appliance. The full-scale DNS service must be explicitly disabled on the member that runs the captive portal. For information on disabling DNS service, see [Starting and Stopping the DNS Service](#).

You can configure one or more captive portals in the Grid. You can also configure one or more member DHCP servers to use a captive portal to register users. For example, if your organization has two sites, you can configure a captive portal for each site and configure the DHCP servers in each site to use their respective captive portals to authenticate users. In order for clients to reach the captive portal, you must specify a route to the captive portal. In a network where all IP addresses are on the same subnet, you can configure Option 33 for the quarantine DHCP range. For additional information, see [Quarantine DHCP Range](#). On a routed network, you must configure a default route on the router for the subnet.

Following are the tasks to configure a captive portal:



1. Select the Grid member that runs the captive portal and configure its properties, as described in the next section, [Configuring Captive Portal Properties](#).
2. Optionally, customize the captive portal and registrati.
3. If you enabled SSL, generate the CA certificate, as described in [Managing Captive Portal Certificates](#) below.
4. Start the captive portal, as described in [Starting the Captive Portal Service](#) below.

## Configuring Captive Portal Properties

When you configure the captive portal properties of a member, you specify if it is used to register users for authentication, guests, or both. If it is used to register guests only, then do not associate it with an authentication server group. You can specify the VIP address of the Grid member or configure an additional IP address on the loopback interface as the captive portal IP address. Alternatively, if the Grid member supports the LAN2 port and it is enabled, but the NIC failover feature is disabled, you can use the IP address of the LAN2 port as the captive portal IP address. To configure an IP address on the loopback interface, see [Configuring IP Addresses on the Loopback Interface](#). For information on the LAN2 port, see [Using the LAN2 Port](#).

In addition, you can configure the port on which the appliance listens for authentication requests redirected from the captive portal. When a user logs in to the captive portal, the member sends an authentication request to its associated authentication server group. The member determines future DHCP replies to client requests based on the authentication result.

To configure the properties of the captive portal:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Captive Portal**.  
Grid Manager lists all the members, except for the Grid Master and Grid Master candidate.
2. Select the member that runs the captive portal and click the Edit icon.
3. In the **General Basic** tab of the *Member Captive Portal Properties* editor, complete the following:
  - **Use This Authentication Server Group for Authenticating Captive Portal Users:** Select the authentication server group that authenticates users for this captive portal. For information about authentication server groups, see [Configuring Authentication Server Groups](#).
  - **Captive Portal User Types:** Specify whether the captive portal is used to register **Authenticated** users only, **Guest** users only, or **Both**.
  - **Portal IP Address:** Select the IP address of the captive portal server. The appliance lists the VIP address and the IP addresses of the loopback interface and the LAN2 port, if enabled. You can select any of these addresses as the portal IP address.
  - **Enable SSL on Portal:** Select this to support encrypted web traffic through SSL/TLS. If you select this option, you must upload a certificate or generate a self-signed certificate. For information about creating and uploading a certificate for the captive portal, see [Managing Captive Portal Certificates](#) below.
  - **Network View:** This field displays if there are multiple network views configured. Select the network view in which the authenticated, quarantine, and guest DHCP ranges belong.
  - **Log Registration Success:** Select to enable the member to log successful registrations in syslog, and then select the logging level from the drop-down list.
  - **Log Registration Failure:** Select to enable the member to log failed registrations in syslog, and then select the logging level from the drop-down list.
4. In the **General Advanced** tab of the editor, you can specify the port on which the member listens for authentication requests redirected from the captive portal. The default port is 4433. Depending on your firewall and network policies, you can configure an unused port greater than 1 and less than 63999.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Customizing the Captive Portal Interface

You can customize the captive portal, and if configured, the guest registration page as well. You can upload image files to the appliance and display your own logo, header and footer. In addition, you can upload the acceptable use policies that are displayed on the captive portal and guest registration page.

Following are guidelines for each item you can customize:

- **Logo Image:** The maximum size is 200 pixels wide by 55 pixels high, and the images can be in JPEG, GIF, or PNG format. It displays on top of the header image.
- **Header Image:** The optimal size is 600 pixels wide by 137 pixels high. The image can be in JPEG, GIF, or PNG format. The header displays at the top of the page.



- **Footer Image:** The optimal size is 600 pixels wide by 20 pixels high. The image can be in JPEG, GIF, or PNG format. The footer displays at the bottom of the page.
- **Acceptable Use Policy:** The policy must be saved as a UTF-8 encoded file. It appears below the welcome message in the captive portal. Users can scroll through the policy when they review it. This is used in the captive portal and guest registrati characters, including white space.

If any of the customizable fields are not configured, then the factory defaults are displayed. To customize the captive portal:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Captive Portal**.
2. Select the member that is running the captive portal and click the Edit icon.
3. Select the **Customization** tab of the *Member Captive Portal Properties* editor.
4. In the **General Captive Portal Customization** section, complete the following:
  - **Company Name:** Enter the name of your company. The company name displays on the title bar of the browser. You can enter a maximum of 256 characters.
  - **Welcome Message:** Type the message that displays on the captive portal. The message can contain a maximum of 300 characters.
  - **Help Desk Message:** Type a message that provides Helpdesk information, such as contact information for technical assistance. The message can contain a maximum of 300 characters.
  - **Logo Image, Header Image, Footer Image, Acceptable Use Policy:** To display the image files and the acceptable use policy on the captive portal, click **Select** beside the item you want to upload. In the *Upload* dialog box, click **Select File** and navigate to the image or text file. Select the file you want to display and click **Upload**. Note that these files have size requirements, as listed earlier in this section.
5. In the **Guest Users Web Page Customization** section, complete the following:
  - The appliance displays certain fields on the guest registration page. Select the checkboxes of the fields that users are required to complete: **Require First Name, Require Middle Name, Require Last Name, Require Email, and Require Phone.**
  - **Custom Field 1 — Custom Field 4:** You can display up to four additional fields on the guest registration page. To add a field to the guest registration page, enter a label for that field. The label can have a maximum of 32 characters. Select **Require** to require users to complete the field. Users can enter a maximum of 128 characters in each of the fields in the captive portal login page and the guest registration page.
6. Save the configuration and click **Restart** if it appears at the top of the screen.

## Managing Captive Portal Certificates

When you enable support for encrypted web traffic sent over SSL/TLS, you can do any of the following:

- Generate a self-signed certificate and save it to the certificate store of your browser.
- Request a CA-signed certificate. When you receive the certificate from the CA, upload it on the member running the captive portal.

## Generating Self-Signed Certificates

You can generate a self-signed certificate for the captive portal. When you generate a self-signed certificate, you can specify the hostname and change the public/private key size, enter valid dates and specify additional information specific to the captive portal. If you have multiple captive portals, you can generate a certificate for each captive portal with the appropriate hostname.

To generate a self-signed certificate:

1. From the Grid tab, select the Grid Manager tab, and then click Captive Portal.
2. Select the member that is running the captive portal, and then click HTTPS Cert -> Generate Self-signed Certificate from the Toolbar.
3. In the Generate Self-signed Certificate dialog box, complete the following:
  - **Secure Hash Algorithm and Key Size:** You can select SHA-1 and a RSA key size of 1024 or 2048. SHA-256 (SHA-2) can be selected together with a RSA key size of 2048 or 4096. The default value is SHA-256 2048.
  - **Days Valid:** Specify the validity period of the certificate.
  - **Common Name:** Specify the domain name of the captive portal.

- **Organization:** Enter the name of your company.
  - **Organizational Unit:** Enter the name of your department.
  - **Locality:** Enter a location, such as the city or town of your company.
  - **State or Province:** Enter the state or province.
  - **Country Code:** Enter the two-letter code that identifies the country, such as US.
  - **Admin E-mail Address:** Enter the email address of the captive portal administrator.
  - **Comment:** Enter additional information about the certificate.
4. Click **OK**.

### Generating Certificate Signing Requests

You can generate a CSR (certificate signing request) that you can use to obtain a signed certificate from your own trusted CA. Once you receive the signed certificate, you can import it in to the Grid member that runs the captive portal, as described in the next section, *Uploading Certificates*.

To generate a CSR:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Captive Portal**.
2. Select the member that is running the Captive Portal, and then click **HTTPS Cert -> Create Signing Request** from the Toolbar.
3. In the *Create Signing Request* dialog box, enter the following:
  - **Secure Hash Algorithm and Key Size:** You can select SHA-1 and a RSA key size of 1024 or 2048. SHA-256 (SHA-2) can be selected together with a RSA key size of 2048 or 4096. The default value is SHA-256 2048.
  - **Common Name:** Specify the domain name of the captive portal.
  - **Organization:** Enter the name of your company.
  - **Organizational Unit:** Enter the name of your department.
  - **Locality:** Enter a location, such as the city or town of your company.
  - **State or Province:** Enter the state or province.
  - **Country Code:** Enter the two-letter code that identifies the country, such as US.
  - **Admin E-mail Address:** Enter the email address of the captive portal administrator.
  - **Comment:** Enter information about the certificate.
4. Click **OK**.

### Uploading Certificates

When you upload a certificate, the NIOS appliance finds the matching CSR and takes the private key associated with the CSR and associates it with the newly uploaded certificate. The appliance then automatically deletes the CSR.

If the CA sends an intermediate certificate that must be installed along with the server certificate, you can upload both certificates to the appliance. The appliance supports the use of intermediate certificates to complete the chain of trust from the server certificate to a trusted root CA.

To upload a certificate:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Captive Portal**.
2. Select the member that is running the captive portal, and then click **HTTPS Cert -> Upload Certificate** from the Toolbar.
3. In the **Upload** dialog box, click **Select File**, navigate to the certificate location, and click **Open**.

The appliance imports the certificate . When you log in to the appliance again, it uses the certificate you imported.

### Downloading Certificates

You can download the current certificate or a self-signed certificate so users can install it in their browsers. To download a certificate:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Captive Portal**.
2. Select the member that is running the captive portal, and then click **HTTPS Cert -> Download Certificate** from the Toolbar.
3. Navigate to where you want to save the certificate and save it.

## Starting the Captive Portal Service

Before you start the captive portal service, ensure that the member is not running any other service. To start the captive portal service:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Captive Portal**.
2. Select the member that is configured to run the captive portal service and click the Start icon.

## Defining IPv4 Network and DHCP Ranges

First define the IPv4 network that uses DHCP authentication, and then define the DHCP ranges and services for each access level that you want to provide on the network:

- Quarantine
- Authenticated
- Guest

For information about configuring DHCP IPv4 networks, ranges and services, see [Managing IPv4 DHCP Data](#), and [Configuring DHCP Properties](#).

### Quarantine DHCP Range

You must configure a DHCP range for the quarantine level so the member DHCP server can assign IP addresses within that range to unauthenticated DHCP clients. An unauthenticated client is allowed to access the captive portal only and must successfully pass the authentication process before it can receive an IP address from the authenticated range. Infoblox recommends 30-second leases for addresses in the quarantine DHCP range. This provides enough time for the user authentication process, so when the client attempts to renew the lease at the midpoint of its lease time, the member can then assign the client a new IP address, depending on the result of the authentication process.

When you configure the quarantine DHCP range, you must specify the captive portal IP address as the DNS server for the address range. The captive portal runs a limited DNS server that resolves all queries with the IP address assigned to the web interface on the captive portal.

Note that you can run the *Captive Portal* wizard to automatically set the lease time of the quarantine range to 30 seconds and to add the captive portal IP address as the DNS server. For information about the *Captive Portal* wizard, see [Using the Captive Portal Wizard](#). Alternatively, you can set the lease time and the DNS server IP address in the **DHCP** tab of the *DHCP Range* editor. For information about the *DHCP Range* editor, see [Configuring IPv4 Address Ranges](#).

To ensure that clients can reach the captive portal, you must specify a route to the captive portal. On a network where all systems can reach each other without going through a router, that is, all IP addresses are on the same subnet, you must configure Option 33 for the quarantine DHCP range. This option specifies a list of static routes that the client should install in its routing cache. The routes consist of a list of IP address pairs. For clients to reach the captive portal, specify the portal IP address first (destination address), and the LAN address of the NIOS appliance second. When the appliance assigns an IP address from the quarantine DHCP range, it also includes the static route that you specified in option 33. For information about configuring DHCP options, see [Configuring IPv4 DHCP Options](#). On a routed network, you must configure a default route via the router on the subnet.

### Authenticated DHCP Range

Configure a DHCP range for authenticated users if you want the Grid member to assign IP addresses within that range to authenticated DHCP clients. Users that receive an IP address in this range typically are allowed full access to the network.

When a client successfully passes authentication, the member automatically stores its MAC address in the corresponding MAC address filter. When the client attempts to renew the lease at the midpoint of its lease time, the member matches the source MAC address in the request with a MAC address in the filter for the authenticated DHCP address range. The member then assigns the client a new IP address from the authenticated DHCP range.

## Guest DHCP Range

Configure a guest DHCP range if you want to provide guest access privileges. You can configure and customize a guest registration page when you configure the captive portal. For information about this feature, see [Customizing the Captive Portal Interface](#).

## Defining MAC Address Filters

After you configure the network and DHCP ranges, you must then configure the MAC address filters and add them to the appropriate DHCP ranges. If you configured DHCP ranges for authenticated and guest users, you must configure MAC address filters for each range with an action of Allow. You must also add those filters to the quarantine range with an action of Deny, to ensure that the member does not allocate an address from the quarantine range to a host whose MAC address matches an entry in the MAC filters for the authenticated and guest DHCP ranges.

When you create the filters, you also specify whether the MAC address entries expire. The member automatically deletes expired MAC address entries from the filter. If a client that registered earlier attempts to renew its IP address or to register after its MAC address has expired, it is redirected to the captive portal because its MAC address is no longer in the filter.

You can run the *Captive Portal* wizard to automatically create the MAC address filters, as described in the next section, [Using the Captive Portal Wizard](#), or you can configure each filter as described in [Defining MAC Address Filters](#).

## Using the Captive Portal Wizard

After you configure the captive portal and the DHCP ranges for each access level, you can use the *Captive Portal* wizard to accomplish the following tasks:

- Associate the captive portal member with the member that serves the DHCP ranges you configured.
- Create MAC address filters and add them to the appropriate DHCP ranges. The wizard allows you to create MAC address filters for the quarantine DHCP range, and for the authenticated and guest DHCP ranges, depending on whether the captive portal is used to register users for authentication, guests, or both. This was specified, when you configured the captive portal properties, described in [Configuring Captive Portal Properties](#). For example, if you indicated that the captive portal is used for authenticated users only, then the wizard allows you to create a MAC filter for the authenticated DHCP range only.
  - If the captive portal is used to register users for authentication, the wizard allows you to create a MAC address filter for the authenticated range. The wizard then automatically adds the filter to the authenticated DHCP range with an action of Allow. It also adds the filter to the quarantine range with an action of Deny. This ensures that the member does not allocate an address from the quarantine range to a host whose MAC address matches an entry in the MAC filter.
  - If the captive portal is used to register guest users, the wizard allows you to create a MAC address filter for the guest range. The wizard then automatically adds the filter to the guest DHCP range with an action of Allow. It also adds the filter to the quarantine range with an action of Deny. This ensures that the member does not allocate an address from the quarantine range to a host whose MAC address matches an entry in the MAC filter.
- Add the captive portal IP address as the DNS server for the quarantine address range.
- Set the lease time of the quarantine range to 30 seconds.

To use the *Captive Portal* wizard to complete the tasks for the DHCP authentication feature:

1. From the **Data Management** tab, select the **DHCP** tab, or from the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and click **Configure Captive Portal**.
3. In the *Captive Portal* wizard, complete the following and click **Next**:
  - **Member DHCP**: Select the member DHCP server that uses this captive portal to authenticate users.
  - **Captive Portal**: Select the member that runs the captive portal. Note that the member that runs the captive portal cannot run any other service, such as DHCP or DNS, and cannot be the Grid Master or Grid Master candidate.
4. This panel allows you to create MAC filters for the authenticated and guest DHCP ranges. The MAC filters you can create depend on your entry in the Captive Portal properties of the Grid member. For example, if you

indicated that the captive portal is used for authenticated users only, then this panel allows you to create a MAC filter for the authenticated DHCP range only.

You can also specify existing MAC filters, if you want to apply them to the authenticated and guest DHCP ranges. Complete the following and click **Next**:

- **Authenticated MAC Filter:** Specify a name for the MAC filter that is used for authenticated users.
  - **Expiration Time:** Specify how long a MAC address is stored in the MAC address filter for authenticated users.
    - **Never:** Select this option to store MAC addresses in the MAC address filter until they are manually removed.
    - **Expires in:** Select this option to store MAC addresses in the MAC address filter for the specified period of time.
  - **Guest MAC Filter:** Specify a name for the MAC filter that is used for guest users.
  - **Expiration Time:** Specify how long a MAC address is stored in the MAC address filter for guest users.
    - **Never:** Select this option to store MAC addresses in the MAC address filter until they are manually removed.
    - **Expires in:** Select this option to store MAC addresses in the MAC address filter for the specified period of time.
5. In this panel, you specify the network and address ranges, so the wizard can apply the MAC address filters to the appropriate ranges. Complete the following:
- **Network:** Select the network that uses DHCP authentication.
  - **Authenticated Range:** Select the IP address range that the appliance uses for authenticated users. The wizard applies the authenticated MAC address filter you specified in the preceding step to this DHCP range with an action of Allow. This effectively allows the member to assign an IP address from the address range to a requesting host whose MAC address matches the MAC address in the filter.
  - **Guest Range:** Select the IP address range that the appliance uses for guest users. The wizard applies the guest MAC address filter you specified in the preceding step to this DHCP range with an action of Allow. This effectively allows the member to assign an IP address from the address range to a requesting host whose MAC address matches the MAC address in the filter.
  - **Quarantine Range:** Select the IP address range that the appliance uses for quarantined addresses. The wizard applies the authenticated and guest MAC address filters to the quarantine DHCP range with an action of Deny. This effectively denies an address request from a host whose MAC address matches an entry in the MAC filters for the authenticated and guest DHCP ranges.
6. Save the configuration and click **Restart** if it appears at the top of the screen.

## Adding and Modifying the Filters and Associations

The *Captive Portal* wizard simplified the configuration process by accomplishing a number of tasks simultaneously. To accomplish each task separately, or to modify the filters or associations after you have run the wizard:

- To define the MAC address filters for each range, see [Defining MAC Address Filters](#).
- To bind each filter to the appropriate DHCP range, see [Applying Filters to DHCP Objects](#).
- To specify the DNS server IP address for the quarantine range and set the lease time to 30 seconds, see [Configuring General IPv4 DHCP Properties](#).
- To associate a member DHCP server with a captive portal and specify the MAC filters for the authenticated and guest DHCP ranges:
  - a. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox -> Edit icon.
  - b. In the *Member DHCP Properties* editor, click the **IPv4 Authenticated DHCP** tab and complete the following:
    - **Use this Captive Portal for Infoblox Authenticated DHCP:** Select this checkbox and select the captive portal that you want to associate with the member.
    - **Authenticated User MAC Filter:** Select the MAC filter used for authenticated users. To change your selection, click **Clear** and click **Select** again.
    - **Guest User MAC Filter:** Select the MAC filter for guest users. To change your selection, click **Clear** and click **Select** again.
  - c. Save the configuration and click **Restart** if it appears at the top of the screen.

## Monitoring DHCP Authentication

You can monitor the status of the captive portal service, as described in [Monitoring Services](#). You can check its status in the *Grid Status* widget and the *Member Status* widget on the Dashboard. For information about these widgets, see [Dashboards](#).

You can also view the MAC addresses that were added to each MAC address filter, as described in [Viewing MAC Address Filter Items](#).

## Viewing DHCP Ranges and Filters

To view the newly created MAC address filters:

1. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab. Grid Manager lists all the configured filters.
2. You can select a filter and view or configure its properties, such as extensible attributes.

For more information about the filters and editing their properties, see [Managing DHCP Filters](#).

To view the DHCP ranges and the newly added filters:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* .
2. Select the DHCP range you want to view and click the Edit icon.
3. If the editor is in Basic mode, click **Toggle Expert Mode**.
4. Click the **Filters** tab to view the filters.

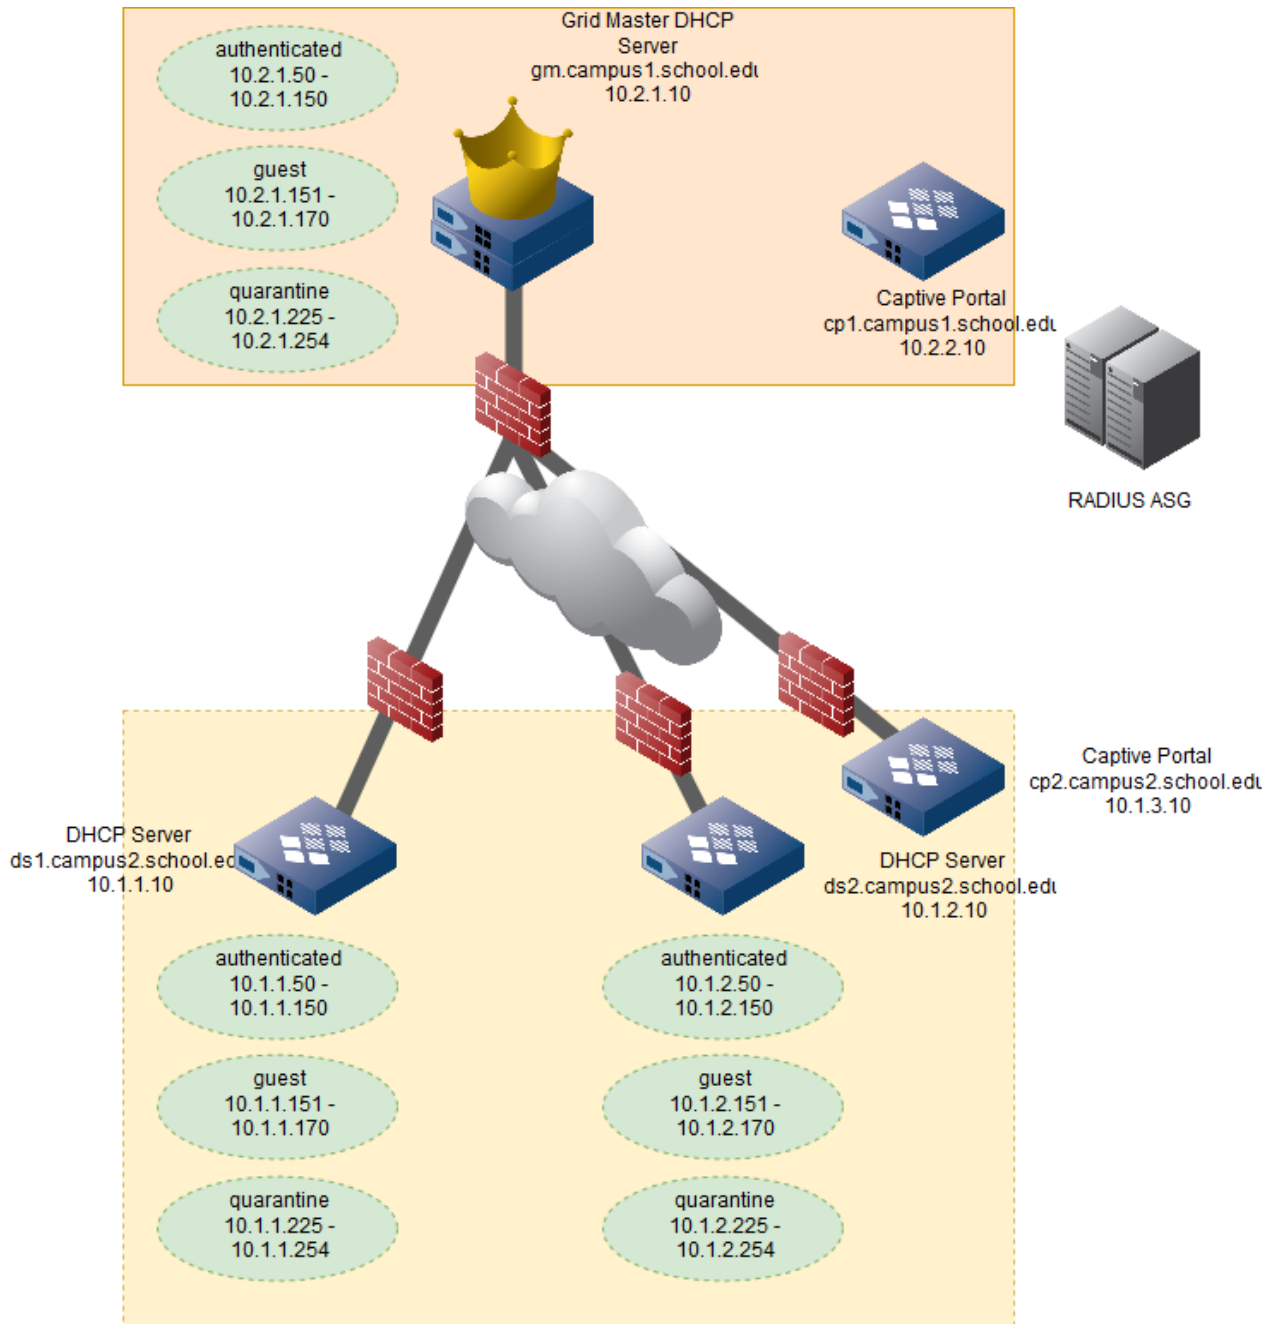
To verify that the captive portal is the DNS server in the quarantine range:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section -> *network* .
2. Select the quarantine DHCP range and click the Edit icon.
3. In the *DHCP Range* editor, click the **DHCP** tab.

The captive portal IP address is listed in the DNS Servers table.

## Configuration Example: Configuring Authenticated DHCP

In this example, a school (school.edu) has two locations, its main campus, campus1.school.edu, and a satellite campus, campus2.school.edu. It has a captive portal server in each location. In the main campus, the Grid Master also functions as a DHCP server and uses a captive portal server to register DHCP clients. In the satellite campus, two members serve DHCP and use the same captive portal server. The captive portal servers use the same RADIUS authentication server group to authenticate users.



## Create the RADIUS Authentication Server Group

Create the RADIUS authentication server group and add two RADIUS servers to the group.

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Expand the Toolbar and click **Add** -> **RADIUS Service**.
3. In the *Add RADIUS Authentication Service* wizard, complete the following:
  - **Name:** Enter **RADIUS ASG**.
  - **RADIUS Servers:** Click the Add icon and enter the following:
    - **Server Name or IP Address:** Enter the RADIUS server FQDN, which is **rs1.school.edu**.
    - **Authentication Port:** Accept the default port (1812).
    - **Authentication Type:** Select the PAP authentication method.



- **Shared Secret:** Enter **no1nose**.
  - **Authentication**
    - **Timeout:** Enter 5 seconds.
    - **Retries:** Accept the default, which is five.
  - **Accounting**
    - **Timeout:** Enter 5 seconds.
    - **Retries:** Accept the default, which is five.
    - Click **Test** to validate the configuration and check that the Grid Master can connect to the RADIUS server.  
Grid Manager displays a message confirming the configuration is valid.
- Click **Add** to add another RADIUS server to the group, and then enter the following:
- **Server Name or IP Address:** Enter the RADIUS server FQDN, which is **rs2.school.edu**.
  - **Authentication Port:** Accept the default port (1812).
  - **Authentication Type:** Select the PAP authentication method.
  - **Shared Secret:** Enter **no1nose**.
- **Authentication**
    - **Timeout:** Enter 5 seconds.
    - **Retries:** Accept the default, which is five.
  - **Accounting**
    - **Timeout:** Enter 5 seconds.
    - **Retries:** Accept the default, which is five.
    - Click **Test** to validate the configuration and check that the Grid Master can connect to the RADIUS server.  
Grid Manager displays a message confirming the configuration is valid.
4. Click **Save & Close**.

## Configure the Captive Portal Properties

Configure the captive portal properties of cp1.campus1.school.edu.

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab -> **Services** tab.
3. Select the member **cp1.campus1.school.edu** and click the Edit icon.
4. In the **General Basic** tab of the *Member Captive Portal Properties* editor, complete the following:
  - **Use This Authentication Server Group for Authenticating Captive Portal Users:** Select **RADIUS ASG**.
  - **Captive Portal User Types:** Select **Both**.
  - **Portal IP Address:** Select **10.2.2.10**.
  - **Enable SSL on Portal:** Select this option.
  - **Log Registration Success:** Select **Informational**.
  - **Log Registration Failure:** Select **Informational**.
5. Click **Save & Close**.

Configure the captive portal properties of cp2.campus2.school.edu.

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab -> **Services** tab.
3. Select the member cp2.campus2.school.edu and click the Edit icon.
4. In the **General Basic** tab of the *Member Captive Portal Properties* editor, complete the following:
  - **Use This Authentication Server Group for Authenticating Captive Portal Users:** Select **RADIUS ASG**.
  - **Captive Portal User Types:** Select **Both**.
  - **Portal IP Address:** Select **10.1.3.10**.
  - **Enable SSL on Portal:** Select this option.
  - **Log Registration Success:** Select **Informational**.
  - **Log Registration Failure:** Select **Informational**.
5. Click **Save & Close**.

## Customize the Captive Portals

Customize the captive portal cp1.campus1.school.edu.

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab-> **Services** tab.
3. Select **cp1.campus1.school.edu** and click the Edit icon.
4. Select the **Customization** tab of the *Member Captive Portal Properties* editor.
5. In the **General Captive Portal Customization** section, complete the following:
  - **Company Name:** Enter **School**.
  - **Welcome Message:** Type the following: **Welcome to School. Please sign in.**
  - **Help Desk Message:** Type: **To reach the Helpdesk, call (408) 111-2222 or email helpdesk@school.edu.**
  - **Logo Image:** Click **Select** beside the logo file and upload it.
6. In the **Guest Users Web Page Customization** section, complete the following:
  - Select the checkboxes beside **Require First Name**, **Require Last Name**, **Require Email**.
7. Click **Save & Close**.

Select the other captive portal server, cp2.campus2.school.edu, and enter the same information.

## Generate a Self-Signed Certificate and Upload It

To generate a self-signed certificate for cp1.campus1.school.edu:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab from the **Services** tab.
3. Select **cp1.campus1.school.edu**, and then click **HTTPS Cert -> Generate Self-signed Certificate** from the Toolbar.
4. In the *Generate Self-signed Certificate* dialog box, complete the following:
  - **Secure Hash Algorithm and Key Size:** You can select SHA-1 and a RSA key size of 1024 or 2048. SHA-256 (SHA-2) can be selected together with a RSA key size of 2048 or 4096. The default value is SHA-256 2048.
  - **Days Valid:** Enter **60 days**.
  - **Common Name:** Enter **cp1.campus1.school.edu**.
5. Click **OK**.
6. Click **Save & Close**.

To generate a self-signed certificate for the captive portal cp2.campus2.school.edu:

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab from the **Services** tab.
3. Select **cp2.campus2.school.edu**, and then click **HTTPS Cert -> Generate Self-signed Certificate** from the Toolbar.
4. In the *Generate Self-signed Certificate* dialog box, complete the following:
  - **Secure Hash Algorithm and Key Size:** You can select SHA-1 and a RSA key size of 1024 or 2048. SHA-256 (SHA-2) can be selected together with a RSA key size of 2048 or 4096. The default value is SHA-256 2048.
  - **Days Valid:** Enter **60 days**.
  - **Common Name:** Enter **cp2.campus2.school.edu**.
5. Click **OK**.
6. Click **Save & Close**.

## Start the Captive Portal Service

1. From the **Grid** tab, select the **Grid Manager** tab.
2. Click the **Captive Portal** tab from the **Services** tab.
3. Select **cp1.campus1.school.edu** and **cp2.campus2.school.edu**, and then click the Start icon.

## Configure the Networks and DHCP Ranges

Configure the network on the Grid Master.

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section.

2. Click the Add drop-down list and select **IPv4 Network**.
3. In the *Add IPv4 Network* wizard, select one of the following and click **Next**:
  - **Add Network**: Click this.
4. Complete the following and click **Next**:
  - **Address**: Enter **10.2.1.0/24**.
5. Complete the following to assign the network to the Grid Master:
  - **Add Infoblox Member**: Select **gm.campus1.school.edu**.
6. Click **Save & Close**.

Configure the ranges on the Grid Master. To create the authenticated range:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** section.
2. Click the 10.2.1.0/24 network link, and then click the Add drop-down list and select **DHCP Range**.
3. In the *Add IPv4 Range* wizard, select **Add DHCP Range** and click **Next**:
4. Complete the following:
  - **Network**: Click **Select Network** and select **10.2.1.0/24**.
  - **Start**: Enter **10.2.1.50**. **End**: Enter **10.2.1.150**.
  - **Name**: Enter **authenticated range**.
5. Click **Next** and complete the following:
  - **Grid Member**: Select this option and select **gm.campus1.school.edu**.
6. Click **Save & Close**.

To create the guest range:

1. Click the 10.2.1.0/24 network link, and then click the Add drop-down list and select **DHCP Range**.
2. In the *Add IPv4 Range* wizard, select **Add DHCP Range** and click **Next**:
3. Complete the following:
  - **Network**: Click **Select Network** and select 10.2.1.0/24.
  - **Start**: Enter **10.2.1.151**.
  - **End**: Enter **10.2.1.170**.
  - **Name**: Enter **guest range**.
4. Click **Next** and complete the following:
  - **Grid Member**: Select this option and select **gm.campus1.school.edu**.
5. Click **Save & Close**.

To create the quarantine range:

1. Click the 10.2.1.0/24 network link, and then click the Add drop-down list and select **DHCP Range**.
2. In the *Add IPv4 Range* wizard, select **Add DHCP Range** and click **Next**:
3. Complete the following:
  - **Network**: Click **Select Network** and select 10.2.1.0/24.
  - **Start**: Enter **10.2.1.225**.
  - **End**: Enter **10.2.1.254**.
  - **Name**: Enter **quarantine range**.
4. Click **Next** and complete the following:
  - **Grid Member**: Select this option and select **gm.campus1.school.edu**.
5. Click **Save & Close**.

Create the network and DHCP ranges for the DHCP servers ds1.campus1.school.edu and ds2.campus2.school.edu.

## Run the Captive Portal Wizard

Run the *Captive Portal* wizard to associate the Grid Master with its captive portal, and to configure the MAC address filters:

1. From the **Data Management** tab, select the **DHCP** tab, or from the **Grid** tab, select the **Grid Manager** tab.
2. Expand the Toolbar and click **Configure Captive Portal**.
3. In the *Captive Portal* wizard, complete the following and click **Next**:
  - **Member DHCP**: Select the Grid Master, **gm.campus1.school.edu**.
  - **Captive Portal**: Select **cp1.campus1.school.edu**.
4. Complete the following and click **Next**:
  - **Authenticated MAC Filter**: Enter **Auth\_MAC\_Filter**.

- **Expiration Time:** Select **Never**.
  - **Guest MAC Filter:** Enter **Guest\_MAC\_Filter**.
  - **Expiration Time:** Select **Never**.
5. Complete the following:
    - **Network:** Select **10.2.1.0/24**.
    - **Authenticated Range:** Select **10.2.1.50 - 10.2.1.150**.
    - **Guest Range:** Select **10.2.1.151 - 10.2.1.170**.
    - **Quarantine Range:** Select **10.2.1.225 - 10.2.1.254**.
  6. Click **Save & Close**.

Run the *Captive Portal* wizard to associate ds1.campus2.school.edu with the captive portal server cp2.campus2.school.edu, and then run it again to associate ds2.campus2.school.edu with the same captive portal server.

## Start the DHCP Service

To start the DHCP service on the Grid Master:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab.
2. Select the Grid Master gm.campus1.school.edu, and the two members, ds1.campus2.school.edu and ds2.campus2.school.edu.
3. Expand the Toolbar and click **Start**.
4. In the *Start Member DHCP Service* dialog box, click **Yes**.
5. Grid Manager starts DHCP services on the Grid Master and on the selected members.

## NAC Integration

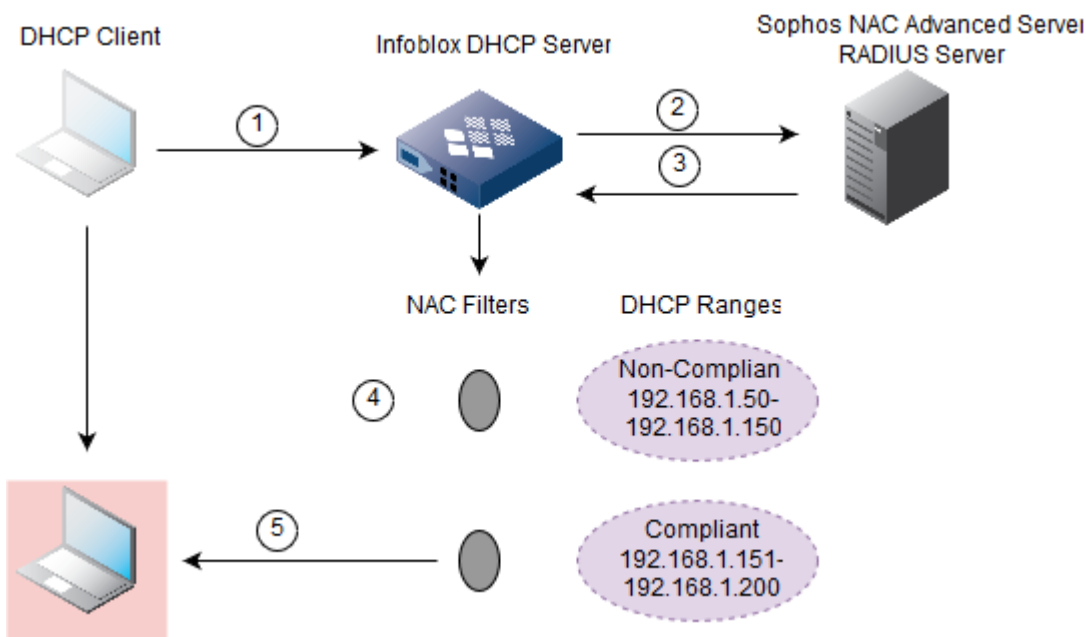
You can configure member DHCP servers to send authentication requests to RADIUS servers and to allocate addresses based on the authentication results. This allows you to place DHCP clients into separate network segments.

You can divide your network into different segments by configuring address ranges and applying NAC filters to them. NAC filters use authentication results from RADIUS servers as matching criteria for granting or denying address requests.

When a DHCP client requests a lease, the member DHCP server can query a remote backend RADIUS server such as the Sophos NAC Advanced server to determine if the DHCP client is authorized to access the network. A Sophos NAC Advanced server is an access-control and compliance server that supports the RADIUS protocol.

The RADIUS server then checks its database and provides the compliance state and user class, if configured, of the DHCP client. The member DHCP server matches the response with the configured NAC filters, and grants a lease to the appropriate network segment.

The following figure presents an example illustrating the authentication process and how a member DHCP server matches the response with NAC filters to determine whether to grant or deny a lease. In the example, there are two DHCP ranges configured, each with a NAC filter that specifies RADIUS compliance state of DHCP clients allowed in each range.



The following steps relate to the above figure.

1. A DHCP client sends a DHCPDISCOVER, DHCPREQUEST, or DHCPINFORM to the Infoblox DHCP server.
2. The DHCP server sends the RADIUS server a RADIUS Access-Request packet that includes RADIUS VSAs (Vendor Specific Attributes) with the MAC address and DHCP transaction ID of the DHCP client.
3. When the RADIUS server receives the Access-Request packet, it does the following:
  - a. It looks up the MAC address in its database to retrieve the associated compliance state and user class.
  - b. The RADIUS server sends back a RADIUS Access-Accept packet that includes RADIUS VSAs with the compliance status and user class.
4. The DHCP server receives the Access-Accept packet and tries to match the response with a NAC filter.
5. The DHCP server matches the response with the NAC filter for compliant DHCP clients and sends the DHCP client a DHCP OFFER that contains an IP address from the corresponding DHCP range. The server also provides the configuration and options associated with that range.

## Configuring NAC with RADIUS Servers

Complete the following tasks to configure the RADIUS server and the member DHCP server. On an already functioning RADIUS server:

- Add the member DHCP server as a RADIUS client. Make sure that the shared secret you enter on the RADIUS server matches the shared secret that you specify when you add the server to the authentication server group in Grid Manager.  
Note that on Grid Manager, you can enter only one shared secret for each RADIUS server. Therefore, on a RADIUS server, you must define the same shared secret for all Grid members that connect to it.  
For information about adding RADIUS clients, refer to the documentation for the RADIUS server.
- Add the Infoblox Grid Master as a RADIUS client, even if it is not going to perform NAC authentication. This enables you to test the connection to the RADIUS server.

On the member DHCP server:

1. Configure the authentication server group for the RADIUS servers. For information, see [Adding a Server Group](#).
2. Associate the authentication server group with the Grid member. For information, see [Associating a Server Group with a Member](#).
3. Configure the network and the DHCP ranges. For information, see [Configuring DHCP Ranges](#).

4. Configure the NAC filters, as described in [About NAC Filters](#).
5. Apply the NAC filters to the DHCP ranges, as described in [Applying Filters to DHCP Objects](#).
6. Enable the DHCP service. For information, see [Starting DHCP Services on a Member](#).

Optionally, you can do the following:

- Manage the authentication cache, as described in [Clearing the Authentication Cache](#).

## About Authentication Servers

You can create a RADIUS authentication server group for Sophos NAC Advanced servers, and then associate the group with the member DHCP server that sends authentication requests. The member DHCP server tries to connect to each Sophos NAC Advanced server in the group using one of the following methods: Ordered List or Round Robin.

In the Ordered List method, the member DHCP server always selects the first Sophos NAC Advanced server in the list when it sends an authentication request. It only sends authentication requests to the next server on the list if the first server goes down.

In the Round Robin method, the member DHCP server selects the first Sophos NAC Advanced server for the first request, the second server for the next request, and so on until it selects the last server in the list. Then it starts with the first server in the list and continues the same selection process.

Each member DHCP server can have only one RADIUS server group assigned, but a RADIUS server group can be assigned to multiple member DHCP servers.

## Adding a Server Group

To create a RADIUS authentication server group for Sophos NAC Advanced servers:

1. From the **Administration** tab, click the **Authentication Server Groups** tab.
2. Expand the Toolbar and click **Add -> RADIUS Service**.
3. In the *Add RADIUS Authentication Service* wizard, complete the following:
  - **Name:** Enter the name of the server group.
  - **RADIUS Servers:** Click the Add icon and enter the following:
    - **Server Name or IP Address:** Enter the Sophos NAC Advanced server FQDN or IP address.
    - **Comment:** You can enter additional information about the server.
    - **Authentication Port:** The destination port on the Sophos NAC Advanced server. The default is 1812.
    - **Authentication Type:** Select the authentication method of the RADIUS server from the drop-down list. You can specify either PAP (Password Authentication Protocol) or CHAP (Challenge Handshake Authentication Protocol). The default is PAP.
    - **Shared Secret:** Enter the shared secret that the member DHCP server and the Sophos NAC Advanced server use to encrypt and decrypt their messages. This shared secret must match the one you entered on the Sophos NAC Advanced server.
    - **Enable Accounting:** Leave this blank. RADIUS accounting is not supported.
    - **Connect through Management Interface:** Select this so that the NIOS appliance uses the MGMT port for communications with just this server.
    - **Disable server:** Select this to disable the Sophos NAC Advanced server if, for example, the connection to the server is down and you want to stop the DHCP server from trying to connect to this server.
    - Click **Test** to validate the configuration and check that the Grid Master can connect to the Sophos NAC Advanced server. Before you can test the configuration though, you must specify the authentication and accounting timeout values.  
If the Grid Master connects to the Sophos NAC Advanced server using the configuration you entered, it displays a message confirming the configuration is valid. If it is unable to connect to the Sophos NAC Advanced server, the appliance displays a message indicating an error in the configuration.
    - Click **Add** to add the Sophos NAC Advanced server to the server group.

When you add multiple Sophos NAC Advanced servers to the list, you can use the up and down arrows to change the position of the servers on the list. The member DHCP server connects to the Sophos NAC Advanced servers in the order they are listed.

- **Authentication**
  - **Timeout:** The time that the member DHCP server waits for a response from a Sophos NAC Advanced server before considering it unreachable. You can enter the time in milliseconds or seconds.
  - **Retries:** The number of times the member DHCP server retries connecting to a Sophos NAC Advanced server before it considers the server unreachable. The default is five.
  - **Mode:** Specifies how the member DHCP server selects the first Sophos NAC Advanced server to contact.
    - **Ordered List:** The member DHCP server always selects the first Sophos NAC Advanced server in the list when it sends an authentication request. It queries the next server only when the first server is considered down. This is the default.
    - **Round Robin:** The member DHCP server selects the first Sophos NAC Advanced server for the first request, the second server for the next request, and so on. If the last server is reached, then the DHCP server starts with the first server in the list, and so on.
  - **Enable Authentication Cache:** The member DHCP server automatically caches authentication results for 120 seconds. When you enable this option, you can override this default in the **Cache Time to Live** field. You must enable this option to clear the cache, as described in [Clearing the Authentication Cache](#) below.
  - **Cache Time to Live:** Specifies the duration of time an authentication result is stored. The default is one hour. The maximum is 259200 seconds (3 days).
  - **Recovery Interval:** Specifies the duration of time a Sophos NAC Advanced server stays inactive after being down, before becoming eligible to have RADIUS requests sent to it. The recovery interval starts when a Sophos RADIUS server is first discovered to be down.
  - **Comment:** You can enter additional information about the server group.
  - **Disable:** Select this to disable the authentication server group.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Associating a Server Group with a Member

To associate an authentication server group with a member DHCP server:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox box, and click the Edit icon.
2. If the *Member DHCP Properties* editor is in Basic mode, click **Toggle Expert Mode**.
3. Select the **IPv4 Authenticated DHCP** tab.
4. Click the **Use this Authentication Server Group for Sophos/RADIUS Authenticated DHCP** checkbox, and then select a group from the drop-down list.
5. Save the configuration and click **Restart** if it appears at the top of the screen.

## Managing Server Groups

To view the list of authentication server groups, from the **Administration** tab, click the **Authentication Server Groups** tab and expand the **RADIUS Service** subtab. For each server group, you can view the server group name, comments, and whether the group is available or disabled. You can then select a server group to modify or delete it.

To modify a server group, select it and click the Edit icon. You can modify any of its properties, and add or delete servers from the group. When you delete a Sophos NAC Advanced server from a group, the appliance permanently deletes it.

To delete a server group, select it and click the Delete icon. When you delete an authentication server group, the appliance permanently deletes it.

## Clearing the Authentication Cache

The authentication cache can store authentication results for up to 20,000 DHCP clients. When the cache reaches its limit, the DHCP member logs a message in syslog. To clear the entire cache or the cache entry of a specific MAC address, you must enable the authentication cache in the RADIUS Service wizard or editor.

To clear the entire authentication cache:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox.
2. Expand the Toolbar and select **Clear** -> **Authentication Cache**.
3. When the **Clear Authentication Cache** confirmation dialog appears, click **Purge**.

To delete a specific entry:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> *member* checkbox.



2. Expand the Toolbar and select **Clear -> Authentication Record**.
3. In the **Clear Authentication Record** dialog box, enter the DHCP client MAC address, and then click **Purge**.

## Configuring DHCP Ranges

Create the IPv4 network and DHCP ranges as described in [Managing IPv4 DHCP Data](#). You can create multiple DHCP ranges and apply one or more NAC filters to each of them.

### Listing DHCP Ranges

By default, DHCP ranges are listed according to their start addresses. You can reorder them according to the order in which you want the member DHCP server to evaluate the ranges.

Consider the following sample DHCP ranges:

- 10.20.30.100-10.20.30.199 (NAC filter that allows leases for compliant DHCP clients)
- 10.20.30.0-10.20.30.99 (No filters)

If the DHCP range with the NAC filter is listed before the range with no filters, then the DHCP server consults the Sophos NAC Advanced server and applies the NAC filter before it grants a lease. It grants leases from the range with no filters only if no NAC filters matched or after all leases from the first range are exhausted. If the first range is the production range and the second range is for the quarantine group, then the server applies the NAC filters for the production range, before it grants leases to the quarantine range.

To change the order of DHCP ranges in a network:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> *network*.
2. Expand the Toolbar and click **Order DHCP Ranges**.
3. In the *Order DHCP Ranges* dialog box, click the up and down arrows to move ranges up or down on the list. The Priority value changes accordingly.
4. Click **OK**.

You can view the DHCP objects in a network, including its DHCP ranges by navigating to the **DHCP** tab -> **Networks** tab -> **Networks** panel, and then clicking the network link. You can select the Priority column for display to view the order of the DHCP ranges. For information about editing the columns, see [Customizing Tables](#).

## About NAC Filters

You can define NAC filters that specify authentication results from a remote, backend RADIUS server such as the Sophos NAC Advanced server. You can then apply each filter to a DHCP range or range template and indicate whether the DHCP server grants or denies a lease when the authentication result matches the filter. You can apply NAC filters to any DHCP range and DHCP range template.

NAC filters are enabled by default. When necessary, you can disable them for the entire Grid so you can perform maintenance on your RADIUS server. When you disable NAC filters, no service interruptions, service down times, configuration changes, or server restarts are required. For information about how to disable NAC filters, see [Disabling NAC Filters](#) below.

In a NAC filter, you can define rules that specify the following:

- The status of the RADIUS authentication server group:
  - Success: At least one of the servers in the RADIUS authentication server group is up.
  - Fail: The MAC address in the DHCP request is not in the authentication cache and all servers in the server group are down.
  - Disabled: The RADIUS authentication server group is disabled, all the servers in the group are disabled, or the member is not assigned a server group.
- The response from the RADIUS server:
  - Accept: The response is an Access-Accept packet.
  - Reject: The response is an Access-Reject packet.
- Whether the Access-Accept packet contains an error. The Infoblox DHCP server expects certain RADIUS VSAs in the Access-Accept packet. An error occurs when any of the RADIUS VSAs are missing. For information about the Access-Accept packet and the RADIUS VSAs, refer to the documentation for the specified RADIUS server.

- Yes: The Access-Accept packet does not include one or more RADIUS VSAs.
- No: There are no errors in the Access-Accept packet.
- A compliance state: unknown, non-compliant, compliant or partially compliant.
- A RADIUS server user class.

When the member DHCP server receives an address request, it checks the DHCP ranges in their priority order. For information about the order of DHCP ranges, see [Listing DHCP Ranges](#).

For each DHCP range, it checks if the request matches any MAC filters, relay agent filters, or DHCP option filters that apply to the range. (For information about these filters, see [Configuring IPv4 DHCP Filters](#).) If any of those filters match, then the member either grants or denies a lease to the DHCP client, based on the filter. If none of those filters match and there are NAC filters defined, then the member tries to send an authentication request to a server in the RADIUS authentication server group.

If you want the member DHCP server to grant leases to specific DHCP ranges in case the RADIUS authentication server group is considered disabled (server state = disabled) or if all RADIUS servers are down (server state = failure), create a NAC filter for each situation and apply it to the appropriate range.

Note that when you create a NAC filter, you do not have to include rules that specify prerequisite conditions. For example, when you create a filter that specifies a RADIUS server compliance state or user class, you do not have to include rules that specify the following: server state=success, server response=accept, and server error = no.

## Defining NAC Filters

To define a NAC filter:

1. From the **Data Management** tab, select the **DHCP** tab -> **Filters** tab, and then expand the Toolbar and click **Add -> IPv4 NAC Filter**.  
or  
From any panel in the **DHCP** tab, expand the Toolbar and click **Add -> IPv4 NAC Filter**.
2. In the *Add Filter Wizard*, complete the following and click **Next**:
  - **Name**: Enter a name for the filter. You can enter a maximum of 255 characters. The name must be unique within a specific network. If you want to specify option settings in the filter, the name must be unique among all NAC filters.
  - **Comment**: Optionally, enter additional information about the NAC filter.
3. Create a rule as follows:
  - In the first drop-down list, select one of the following criterion: **Compliance State**, **Server Error**, **Server Response**, **Server State** or **User Class**.
  - In the second drop-down list, select an operator: **equals** or **does not equal**.
  - The selections in the third drop-down list depend on the criterion you selected:
    - **Compliance State**: Select one of the following compliance states: **Unknown**, **Non-compliant**, **Compliant** or **Partially Compliant**.
    - **Server Error**: The Infoblox DHCP server expects certain RADIUS VSAs in the Access-Accept packet. When any of the VSAs are missing, then the DHCP server considers this an error. For information about the Access-Accept packet and the VSAs, refer to the documentation for the specified RADIUS server. Select one of the following:
      - **Yes**: Create a rule that matches when the RADIUS server sends an Access-Accept packet with a missing VSA.
      - **No**: Create a rule that matches when the RADIUS server sends an Access-Accept packet with no errors.
    - **Server Response**: Select one of the following:
      - **Accept**: Create a rule that matches when the server sends back an Access-Accept packet.
      - **Reject**: Create a rule that matches when the server sends back an Access-Reject packet.
    - **Server State**: Select one of the following:
      - **Success**: Create a rule that matches when at least one RADIUS server in the group is up.
      - **Fail**: Create a rule that matches when the MAC address of the DHCP client is not in the cache and all RADIUS servers in the server group are down.
      - **Disable**: Create a rule that matches when the RADIUS authentication server group is disabled, all servers in the group are disable, or the member was not assigned a server group.

**User Class:** Enter the RADIUS user class value, for example, NACDeny. The member DHCP server does not validate the entry. Therefore, you must make sure that the user class you enter matches the user class name on the RADIUS server.

To add another rule:

- Click **+** to add another rule at the same level.
- Click **|<** to add an **all** (logical AND) or **any** (logical OR) operator line and a parenthetical rule that is indented one level and above the first rule.
- Click **->** to add an **all** (logical AND) or **any** (logical OR) operator line and a parenthetical rule that is indented one level.

After you add all the match rules, you can click **Preview** to view the rules or click **Reset** to remove the previously configured rules and start again.

4. Click **Next** and complete the following to define DHCP options:

- **Option Space:** Select an option space from the drop-down list. This field is not displayed if you do not have custom option spaces. The appliance uses the **DHCP** option space as the default.
- **Lease Time:** Enter the value of the lease time in the field and select the time unit from the drop-down list. The lease time applies to hosts that meet the filter criteria.

#### Options to Merge with Object Options

Click the Add icon. Grid Manager adds a new row to the table with the default **DHCP** option space and option name displayed. Complete the following:

- **Option Space:** Click the down arrow and select an option space from the drop-down list. The selected option space contains the corresponding DHCP options.
- **Option Name:** Click the down arrow and from the drop-down list, select the DHCP option you want to return to the matching client.
- **Value:** Enter the match value that you want the filter to use for the selected DHCP option. For example, enter the value 172.124.3.0 for the SUNW.SrootIP4 option.

To add more options to the filter, click the Add icon and repeat the steps.

5. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).
6. Save the configuration and click **Restart** if it appears at the top of the screen.

After you add NAC filters, you must then apply them to DHCP ranges, as described in [Applying Filters to DHCP Objects](#). You can also list, modify or delete NAC filters, as described in [Managing DHCP Filters](#).

## Disabling NAC Filters

NAC filters are enabled by default. When you disable them, the appliance bypasses evaluations of all NAC filters for the entire Grid. There are no configuration changes, service restarts, or service down times when you disable the NAC filters. The appliance keeps the filter configurations so you can enable them at a later time.

To disable NAC filters for the Grid:

1. From the **Data Management** tab -> **DHCP** tab, select **Grid DHCP Properties** from the Toolbar.
2. In the *Grid DHCP Properties* editor, click **Toggle Advanced Mode**, select the **General** tab -> **Advanced** tab, and then complete the following in the Common Properties section:
  - **Disable All NAC Filters:** Select this to disable all NAC filters in the Grid. The appliance keeps the filter configurations so you can enable them when needed.

## Managing Leases

Historical DHCP lease records complement the real-time DHCP lease viewer by allowing the appliance to store and correlate DHCP lease information over the lifetime of a lease. You can see critical information such as when the appliance issued or freed an IPv4 or IPv6 address, the MAC address or DUID and host name of the device that received the IP address, the Grid member that supplied the lease, and the start and end dates of the lease.

You can view current leases and lease history in the **Data Management** -> **DHCP** -> **Leases** tab in Grid Manager. To view lease history, you must first enable lease logging at the Grid level. For information, see [Configuring DHCP Logging](#) and [Configuring the Lease Logging Member](#). You can also export the DHCP lease history log in CSV format for archival and reporting purposes.

In the **Leases** tab, you can do the following:

- View current leases. For information, see [Viewing Current Leases](#).
- View detailed information about a specific lease. For information, see [Viewing Detailed Lease Information](#).
- View historical lease records. For information, [Viewing Lease History](#).
- Export current leases and lease history logs. For information, see [Exporting Lease Records](#).
- Clear leases. For information, see [Clearing Leases](#).

You can also use the filter and **Go to** functions in the lease panels to retrieve lease information for specific hosts, MAC addresses, and IP addresses. These capabilities are crucial for security auditing and for meeting new compliance regulations such as SOX and HIPAA. You can also sort the lease information by column.

This section explains how to manage IPv4 and IPv6 leases. It contains the following topics:

- [Viewing Current Leases](#)
- [Viewing Detailed Lease Information](#)
- [Viewing Lease History](#)
- [Exporting Lease Records](#)
- [Clearing Leases](#)

## Viewing Current Leases

To view current IPv4 and IPv6 leases:

1. From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Current Leases**.
2. Grid Manager displays the following information:
  - **IP Address:** The IPv4 address or IPv6 prefix or address that the appliance assigned to a DHCP client for this lease.
  - **Protocol:** Indicates whether the lease is for an IPv4 or IPv6 address.
  - **Members/Servers:** The Grid member or Microsoft server (for IPv4 leases only) that granted the lease.
  - **MAC address:** The MAC address of the IPv4 DHCP client that received the lease for an IPv4 address.
  - **DUID:** The DHCP Unique Identifier (DUID) of the IPv6 DHCP client that received the lease for an IPv6 address.
  - **Host Name:** The hostname that the DHCP client sent with its DHCP request. For IPv4 leases, this field displays the hostname of the DHCP client. For IPv6 leases, this field typically displays the FQDN.
  - **Fingerprint:** The name of the DHCP fingerprint or vendor ID of the DHCP client that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#).
  - **State:** The binding state of the current lease. The lease state can be one of the following:
    - **Free:** The lease is available for clients to use.
    - **Active:** The lease is currently in use by a DHCP client.
    - **Static:** The lease is a fixed address lease.
    - **Expired:** The lease was in use, but the DHCP client never renewed it, so it is no longer valid.
    - **Released:** The DHCP client returned the lease to the appliance.
    - **Abandoned:** The appliance cannot lease this IP address because the appliance received a response when pinging the address.
  - **End:** The day, date, and time when the state of the lease ends.
  - **Start:** The day, date, and time when the state of the lease starts.
  - **Username:** Displays the name of the user who receives the lease for the IP address. The username enables you to differentiate between guest users and authenticated users. If you log in as an authenticated user, your username is whatever you choose when you log in. If you log in as a guest, your username is First: first\_name Last: last\_name.  
For example, if your first name is John and last name is Doe and your username is jdoe, when you log in

as an authenticated user, your username is jdoe. If you log in as a guest user, your username is First: John, Last: Doe.

- **Client ID:** The DHCP client identifier (option 61) in an IPv4 lease. The client sends the client identifier as option 61 in the DHCP DISCOVER and REQUEST packets, as described in *RFC2132, DHCP Options and BOOTP Vendor Extensions*. The client identifier is either the MAC address of the network interface card requesting the address or any string uniquely identifying the client. This field is not displayed by default.

Note that the dates and timestamps in the Leases tab are determined by the time zone setting of the admin account that you use to log in to the appliance.

You can display the following discovered data for IPv4 leases:

- **Last Discovered:** The timestamp when the IP address was last discovered. This data is read-only.
- **OS:** The operating system of the detected host or virtual entity. The OS can be one of the following:
  - **Microsoft** for all discovered hosts that have a non-null value in the MAC addresses using the NetBIOS discovery method.
  - A value that a TCP discovery returns.
  - The OS of a virtual entity on a vSphere server.
- **NetBIOS Name:** The name returned in the NetBIOS reply or the name you manually register for the discovered host.
- **Discovered Name:** The name of the network device associated with the discovered IP address.
- **Discoverer:** Specifies whether the IP address was discovered by a PortIQ or NIOS discovery process.

You can do the following in this tab:

- Sort the data in ascending or descending order by column.
- View the lease detailed information of a current lease by selecting the checkbox of the lease, and then clicking the Open icon.
- Change a current lease state to **Free** by selecting the checkbox of a current lease, and then clicking the Delete icon.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Print and export the data in this tab.

## Viewing Detailed Lease Information

You can view detailed information about a specific lease. To view detailed information of a specific lease:

1. From the **Data Management** tab, select **DHCP** tab -> **Leases** -> **Current Leases** -> *lease* checkbox, and then click the **Lease Details** icon.  
or  
From the **Data Management** tab, select the **IPAM** tab, drill down to the IP Map, IP List, or IP address panel, and then click **Lease Details** from the Toolbar.
2. In the *Lease Detailed Information* viewer, Grid Manager displays the following for each type of lease:  
For IPv4 leases, it displays the fields **Member**, **MAC address**, **Host**, **Start**, **End**, **Binding State**, **Username**, **Binding State**, as described in [Viewing Current Leases](#), plus the following information:
  - **Lease Issue:** The date and time when the lease was issued. Displayed in the lease event details report only.
  - **Event:** The action taken. This can be one of the following: **Issued**, **Renewed**, **Freed**, or **Abandoned**. Displayed in the lease event details report only.
  - **Served by:** The member that provides DHCP services to the lease.
  - **Next Binding State:** The subsequent binding state when the current lease expires. The lease state and the next binding state can be one of the following:
    - **Free:** The lease is available for clients to use.
    - **Active:** The lease is currently in use by a DHCP client.
    - **Static:** The lease is a fixed address lease.
    - **Expired:** The lease was in use, but the DHCP client never renewed it, so it is no longer valid.
    - **Released:** The DHCP client returned the lease to the appliance.

- **Abandoned:** The appliance cannot lease this IP address because the appliance received a response when pinging the address.
- **Billing Class:** The billing class of the lease.
- **Option 82 Agent ID:** The agent ID of the relay agent filter (option 82). A relay agent can append DHCP option 82, relay agent information, to a message that it forwards from a DHCP client to a DHCP server.
- **Option 82 Circuit ID:** The circuit ID of the relay agent filter (option 82).
- **Option 82 Remote ID:** The remote ID of the relay agent filter (option 82).
- **Option 82 Link Selection:** An IP address, provided by the DHCP relay agent, in the subnet from which the DHCP server allocated the leased IP address to the client.
- **Option 82 Server ID Override:** The IP address, provided by the DHCP relay agent, that the DHCP server used as the value in the server identifier option of the response sent to the client.

For more information about the sub options for Option 82, see [About the DHCP Relay Agent Option \(Option 82\)](#).

The agent, circuit, and remote IDs for option 82 can be displayed in hexadecimal or plain text format. By default, Grid Manager displays them in hexadecimal format. You can change the logging format, as described in [Defining Logging Format for DHCP Option 82](#).

- **Option:** Agent circuit ID and remote ID data sent by a DHCP relay agent in all DHCP options.
- **UID:** (User ID) The client identifier that the DHCP client sends the appliance (in DHCP option 61) when it acquires the lease. Not all DHCP clients send a UID.
- **TSFP:** (Time Sent From Partner) The time — from the point of view of a remote DHCP failover peer — when the current lease state ends.
- **CLTT:** (Client Last Transaction Time) The time of the last transaction with the DHCP client for this lease.
- **TSTP:** (Time Sent To Partner) The time — from the point of view of the local DHCP failover peer — that the current lease state ends.

For IPv6 leases, it displays most of the same fields as IPv4 leases, plus the following information:

- **DUID:** The DUID of the IPv6 DHCP client that received the lease for an IPv6 address.
- **Prefix Bits:** The prefix length.
- **Preferred Lifetime:** The length of time that a valid address is preferred. A preferred address can be used with no restrictions. When this time expires, the address becomes deprecated.

## Viewing Lease History

To view lease history:

- From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Lease History**.

Grid Manager displays a table of historical leases that have been archived in the system. You can export the information in the lease history table. You can also search by the IP address or MAC address of the lease. Grid Manager displays the following read-only information:

- **Lease Issue:** The date and time when the lease was issued.
- **Protocol:** Indicates whether the lease is for an IPv4 or IPv6 address.
- **IP Address:** The IPv4 address or IPv6 prefix or address of the lease.
- **MAC Address:** The MAC address of the IPv4 lease.
- **DUID:** The DUID of the DHCP client that received the lease for an IPv6 address.
- **Fingerprint:** The name of the DHCP fingerprint or vendor ID of the DHCP client that was identified through DHCP fingerprint detection. This field displays **No Match** for devices that do not have any DHCP fingerprint information. For information about DHCP fingerprints, see [DHCP Fingerprint Detection](#).
- **Host Name:** The host name that the DHCP client sent to the appliance.
- **Action:** This can be one of the following: Issued, Renewed, Freed, or Abandoned.
- **User Name:** The name of the user who received the lease for the IP address.
- **Start:** The start date of the lease.
- **Stop:** The end date of the lease.
- **Member/Server:** The DHCP member or Microsoft server that granted the lease.
- **Member IP Address:** The IP address of the DHCP member that granted the lease. You can do the following in this section:



- View the lease event detailed information of a historical lease by selecting the checkbox of a lease, and then clicking the Open icon.
- Print or export the information in this section.

## Viewing Lease Event Detailed Information

You can view detailed information about a historical lease record by clicking the lease in the **Data Management** tab -> **DHCP** tab -> **Leases** tab -> **Lease History**. Grid Manager displays the event, the date and time when the event occurred, plus detailed information about the historical lease record. For information about the fields, see [Viewing Detailed Lease Information](#).

You can also export and print the information in this panel. For information, see [Exporting Lease Records](#).

## Exporting Lease Records

The DHCP lease history log holds a maximum of 100,000 entries. After that maximum is reached, the appliance begins deleting entries, starting with the oldest. To archive DHCP lease history logs, you can export them and save them as CSV (comma separated variables) files. You do not need to export the entire log. You can selectively export a section of the log, such as the lease events for a single day.

As a conservative approach to archiving DHCP lease data, Infoblox recommends exporting the log on a daily basis, perhaps through API (application programming interface) scripting. By exporting the daily log entries every day over a certain period of time and then opening the exported files with a spreadsheet program, you can see the number of entries for each day. You can then estimate how often you need to export the log to ensure that you save all of the entries before the log fills up (at 100,000 entries). As a result, you might discover that you need to export the log more or less frequently than once a day to archive all the records.

A limited-access admin group can view and export the DHCP lease history if it has read-only permission to the DHCP lease history. For information on setting permissions for the DHCP lease history, see [Administrative Permissions for the IPv4 and IPv6 DHCP Lease Histories](#). In addition, you can export the displayed DHCP current lease information or you can export them to a CSV file.

To export displayed current lease information:

1. From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Current Leases**.
2. Click the Export icon and select **Export visible data**. For more information on how to export, see [Exporting Displayed Data](#).

To export DHCP current lease information to a CSV file:

1. From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Current Leases**.
2. Click the Export icon and select **Export data in Infoblox CSV Import format**. For more information on how to export, see [Exporting Data to Files](#).

To export a lease history log:

1. From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Current Leases** or **Lease History**.
2. Click the Export icon and select.
3. In the *Export* dialog box, click **Start**.
4. Click **Download** when the export is complete. Ensure that you turn off the pop-up blocker in your browser.
5. In the *File Download* dialog box, select the appropriate action to either open or save the CSV file.

## Clearing Leases

You can clear active leases for which you have read/write permission. When you clear an active lease, its IP address becomes available and its status changes to "Free". To clear an active lease:

1. From the **Data Management** tab, select the **DHCP** tab -> **Leases** tab -> **Current Leases**.
2. Click the checkboxes beside the IP addresses of the leases you want to clear, and then click the Clear Lease icon.



Grid Manager clears the selected leases. You can view information about a cleared lease, by selecting it in the Lease History panel and clicking the Edit icon.

## Configuring Microsoft Windows Servers

This section describes how you can centrally manage Microsoft Windows® DNS and DHCP servers from Grid Manager. You can synchronize your DNS and DHCP data from the Microsoft servers to the Grid, and then use IPAM tools to facilitate DHCP and DNS configuration and data management. This section includes the following topics:

- [Managing Microsoft Windows Servers](#)
- [Managing Microsoft DNS Services](#)
- [Managing Microsoft DHCP Services](#)

## Managing Microsoft Windows Servers

This section explains how to configure Grid members to manage Microsoft Windows DNS and DHCP servers from Grid Manager. It includes the following topics:

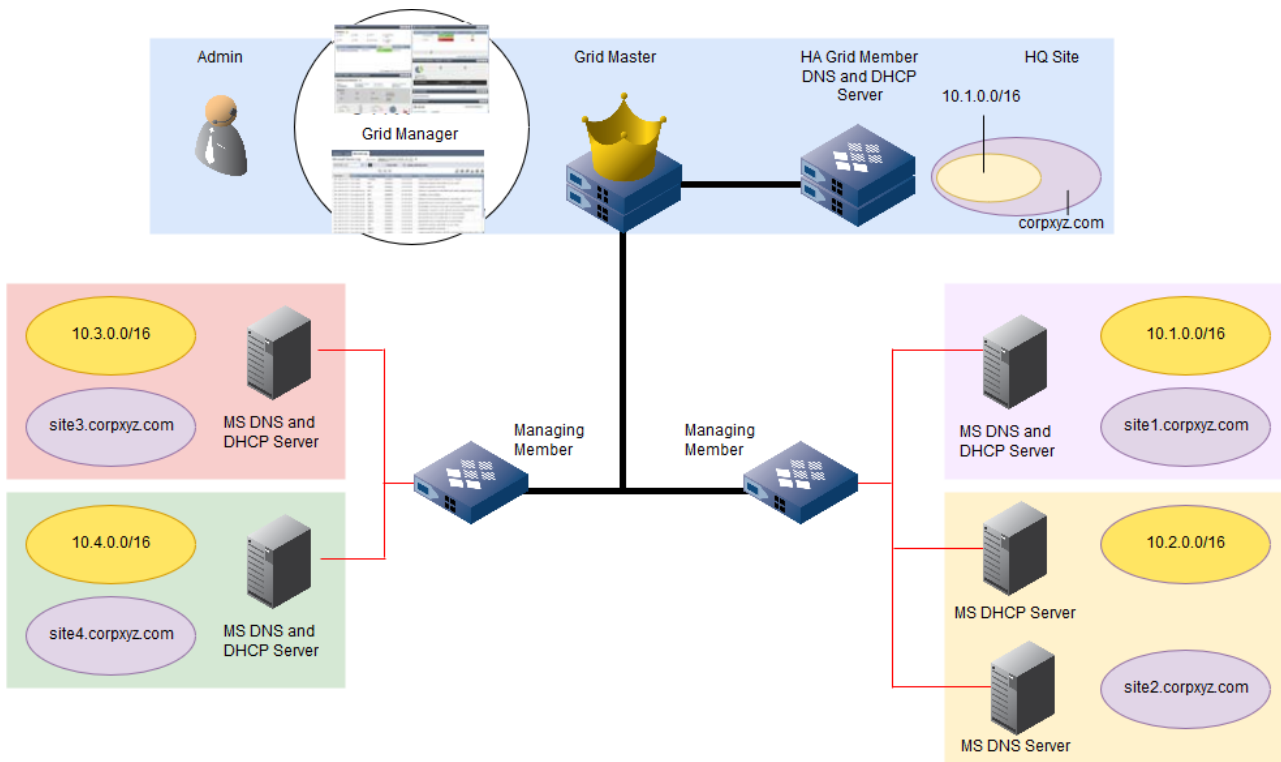
- [About Managing Microsoft Windows Servers](#)
- [Configuring Members to Manage Microsoft Servers](#)
- [Configuring Grid Properties for Managing Microsoft Servers](#)
- [Managing Microsoft Servers](#)
- [Configuring Active Directory Sites and Services](#)
- [Configuring Active Directory Sites and Associated Networks](#)
- [Managing Active Directory Sites](#)
- [Managing Active Directory Sites and Associated Networks](#)
- [Configuring Identity Mapping](#)
- [Monitoring Managed Microsoft Servers](#)

## About Managing Microsoft Windows Servers

You can configure Grid members to manage Microsoft Windows DNS and DHCP servers, and synchronize their DNS and DHCP data to the Grid database, so you can view and optionally, manage the data from Grid Manager. After the data is synchronized, you can use the IPAM tools of Grid Manager to simplify DNS and DHCP configuration and troubleshooting. You can use Smart Folders to organize your data, and monitor your networks and Microsoft servers from the Dashboard. In addition, you can control the DNS and DHCP services of the Microsoft servers from Grid Manager and configure server properties as well. You can use the Identity Mapping feature to get visibility of user interaction with their environments. For more information about the Identity Mapping feature, see [Configuring Identity Mapping](#).

The below figure illustrates a Grid that includes a member that provides DNS and DHCP services, and two other members that manage multiple Microsoft DNS and DHCP servers. Assuming the admin has the appropriate permissions, the admin can centrally manage the Microsoft DNS and DHCP servers and Infoblox DNS and DHCP server from a single interface, Grid Manager.

*Managing Microsoft and Infoblox DNS and DHCP Servers from the Grid Master*



You do not have to configure or install any application on the Microsoft servers for the Grid members to communicate with the servers. Infoblox uses MS-RPC (Microsoft Remote Procedure Calls) to manage Microsoft servers.

A Grid member can manage a Microsoft server in either of two modes, Read-only or Read/Write. In Read-only mode, the Grid member synchronizes data from the Microsoft server to the Grid so admins can use Grid Manager to view the synchronized data, but not update it. Read/Write mode allows admins to update the synchronized data as well.

Updates from Grid Manager are then synchronized to the Microsoft server, and updates from the Microsoft server are synchronized to the Grid.

Configuration changes and data synchronized from the Grid to the Microsoft server apply immediately after the synchronization. You do not have to restart the Microsoft server or for DNS, reload the zones.

Note that due to a field length limit set on the Microsoft DHCP server, after you synchronize DHCP data on the Microsoft server, the "Comment" and "Description" fields for a fixed address and reservation can display only up to 128 characters even though NIOS allows up to 256 characters for these fields.

**Note**

A Grid member must have a Microsoft Management license installed to manage a Microsoft server. The license allows the member to synchronize data with Microsoft servers. It also activates the tabs, dialog boxes and other elements in Grid Manager that you need to manage a Microsoft server. If you do not see the Microsoft Servers tab after you add a member that has a Microsoft Management license, you might have to restart the Grid Master to view the tab and to manage Microsoft DNS and DHCP servers in the Grid.

### Supported Windows Versions

Infoblox Grid members can manage Microsoft servers that support the following Windows versions:

OS	Levels	Platforms
Microsoft Windows 2003 Standard and Datacenter	SP2	32 bits
Microsoft Windows 2003 R2 Standard and Datacenter	Initial Release	32 bits, 64 bits
Microsoft Windows 2008 Standard and Datacenter	SP2	32 bits, 64 bits
Microsoft Windows 2008 R2 Standard and Datacenter	Initial Release	64 bits
Microsoft Windows 2012 Standard and Datacenter	Initial Release	64 bits
Microsoft Windows 2012 R2 Standard and Datacenter	Initial Release	64 bits
Microsoft Windows 2016 Standard and Datacenter	Initial Release	64 bits
Microsoft Windows 2019 Standard and Datacenter	Initial Release	64 bits

Infoblox supports the following SMB (Server Message Block) protocol versions for Microsoft Windows servers: SMB version 1 (SMBv1), SMB version 2.x (SMBv2.x), and SMB version 3.x (SMBv3.x).

Grid members check the Windows version of the Microsoft servers before each synchronization. If a Microsoft server reports an unsupported version before a synchronization, the member logs an error and the synchronization fails. Note that some Windows versions require certain updates and hotfixes installed, so the Microsoft server can synchronize with the Grid member. Following are the current requirements:

- Windows Server 2003, Enterprise x64 Edition requires the installation of security update 935966.
- Windows Server 2008 R2 requires the hotfix referenced in the Knowledge Base article 981776.
- Windows Server 2008-based DNS servers might not display delegations for reverse lookup zones. For information about this issue, including the available hotfix, refer to Knowledge Base article 958190.

For information about the updates, enter their IDs in the Search field of the Microsoft Support website at <http://support.microsoft.com>.

## Administrative Permissions

By default, only superusers can configure Grid members to manage Microsoft servers. Superusers can give limited-access users Read-only or Read/Write permission to Microsoft servers. Read-only permission allows admins to view the properties and data of a Microsoft server from Grid Manager. Write permission is required to configure Grid members to manage Microsoft servers, edit their properties, and start or stop their DNS and DHCP services. For additional information, see [Administrative Permissions for Microsoft Servers](#).

Note that to view and manage the DNS and DHCP data synchronized from Microsoft servers, admins must have permissions to the applicable DNS and DHCP resources. For example, to view DNS zones synchronized from Microsoft servers, admins must have Read-only permission to zones, and to edit the zones, admins need Read/Write permission to them. Similarly, to view DHCP ranges synchronized from Microsoft servers, admins must have Read-only permission to DHCP ranges, and to edit the DHCP ranges, admins need Read/Write permission to the DHCP ranges. For information, see [Administrative Permissions for DNS Resources](#) and [Administrative Permissions for DHCP Resources](#).

The administrative permissions on the Grid are different from those on the Microsoft server. These permissions are independent of each other and are not synchronized.

## Deployment Guidelines

Following are some recommendations and considerations when configuring Grid members to manage Microsoft servers:

- Infoblox recommends that you schedule the initial synchronization at a time when your network is less busy, especially if you are synchronizing a large amount of data. In addition, if a Microsoft server reconnects after being disconnected for a long period of time, it could synchronize a significant amount of data and this could impact the Grid Master performance.
- vNIOS Grid members and Grid members running on Infoblox-250, TrinziC 100, and TrinziC 810 appliances do not support being configured as managing members.
- The managing member must be close, in terms of network hops, latency and bandwidth, to the Microsoft servers that it manages. This will help reduce the synchronization time and potential retries due to network delays.
- Although a Grid member that manages Microsoft servers can run other protocols and services, to optimize performance, Infoblox recommends that you configure one or more members solely for managing Microsoft servers.
- Grid members connect to Microsoft servers using RPC calls over TCP/IP. You must adjust your firewall policies to allow traffic between the managing Grid member and its assigned Microsoft servers. Grid members use the VIP as their source port. In Windows Server 2003, RPC uses the dynamic port range 1025-5000, by default. In Windows Server 2008, RPC uses the dynamic port range 49152-65535, by default. You can reduce the number of available ports as follows:
  - In Windows Server 2003, use the rpccfg.exe tool. For information, refer to <http://support.microsoft.com/kb/908472>.
  - In Windows Server 2008 and later, use the netsh tool. For information, refer to <http://support.microsoft.com/kb/929851>.

The minimum number of ports required in the range is 255.

Note that TCP ports 135 and 445 must be open on the Microsoft server, in addition to the dynamic port range. Ports 135 and 445 are used by the port mapper interface, which is a service on the Microsoft server that provides information to clients on which port to use to connect to a specific service, such as the service that allows the management of the DNS service.

- The capacity of the managing member must be greater than or equal to the sum of all its assigned Microsoft servers.
- The capacity of the Grid Master must be greater than or equal to the sum of all managed Microsoft servers
- A Microsoft server can synchronize its data to only one network view, and for DNS data, only one DNS view.
- Multiple Microsoft servers can synchronize their data into the same network view and DNS view, unless there is a conflict in their data. For example, two Microsoft servers in different locations could serve the same private IP address space, such as 10.1.0.0/16, or serve reverse-mapping zones with the same name, such as 10.in-addr.arpa. Synchronizing their data to the same network view and DNS view would cause conflicts which result in the Grid member synchronizing the data of only one Microsoft server and logging an error for the other Microsoft server. In such situations, Infoblox recommends that you synchronize each Microsoft server to a different network view and DNS view to ensure that data from both servers are synchronized.
- This release supports the following Microsoft IPAM enhancements:
  - Monitor and control settings for DNS and DHCP services for Microsoft servers
  - Synchronization of IP addresses with invalid MAC addresses
  - Output destination for Microsoft server log messages in the syslog
  - Synchronization and configuration of Microsoft DHCP failover relationships
  - RPC (Remote Policy Call) timeout setting
  - Maximum concurrent connections for Microsoft servers
  - Enabling and Disabling DNS zone synchronization
  - The ability to allow GSS-TSIG based DDNS updates from multiple clients in a single forest or multiple forests using keys that are appropriate for their respective domains.

Earlier NIOS releases do not support these features. When you modify settings related to these features on Microsoft servers assigned to Grid members running a NIOS release earlier than NIOS 6.11, the NIOS appliance displays an error message.

## Limitations for Scheduling Full Upgrades

When you schedule a full upgrade to NIOS 6.11 and later, you can set the RPC (Remote Procedure Calls) timeout settings and maximum concurrent connections immediately on any member as soon as you upgrade the Grid Master. When you disable the synchronization of a DNS zone, NIOS displays an error message if any Microsoft server that is a potential synchronization master, which is assigned to a Grid member, has not been upgraded to NIOS 6.11.0 or later

versions.

The following happens if a DHCP failover configuration exists:

- Any Microsoft server configured to manage DHCP and assigned to a Grid member, which has not been upgraded, synchronizes the DHCP configuration without any failover related data.
- You cannot assign a Microsoft server that is assigned to a Grid member, which has not been upgraded, to any Microsoft failover relationship.
- NIOS manages any DHCP scope that is identified to be part of a Microsoft failover relationship but is managed by a Grid member that has not been upgraded, in Read-only mode.

## Configuring Members to Manage Microsoft Servers

You can manage Microsoft DNS and DHCP servers on any Grid member. To avoid performance issues, Infoblox strongly recommends that you do not configure Microsoft DNS and DHCP servers on the Grid Master and Grid Master candidate. When an HA pair manages Microsoft servers, the active node handles synchronization. If an HA failover occurs during a synchronization, the failing node immediately aborts the synchronization. The new active node resumes the next synchronization. Changes that occurred on the Grid since the end of the last synchronization are lost.

For Microsoft DHCP failover, NIOS supports both the hot standby and load sharing modes in both Read/Write and Read-only modes on DHCP servers running Microsoft Windows 2012 and 2012 R2. For more information about Microsoft DHCP failover, refer to the Microsoft documentation.

Complete the following tasks to configure a Grid member to manage a Microsoft server:

1. On the Microsoft server, create a user account for the Grid member. For information, see [Setting Microsoft Server Credentials](#) below.
2. On the Grid Master, configure the managing member, as described in [Configuring a Managing Member](#) below.

### Setting Microsoft Server Credentials

To enable a Grid member to synchronize data with a Microsoft server and control DNS and DHCP services, you must do the following on the Microsoft server:

1. Create a user account for the Grid member.
2. Grant the user account the necessary permissions.

You can either add the user account to the Administrators Group or add the user account to specific groups and explicitly set only the permissions necessary to access the DHCP and DNS services of the Microsoft server. The following sections provide general instruction on each method.

#### Adding User Account to the Administrators Group

Adding the user account of the Grid member to the Administrators Group provides total control over the Active Directory Domain. Do one of the following:

- If the managed Microsoft server is a standalone server or a member server in a domain, open **Computer Management**, click **Groups**, and add the user account to the Administrators Group.
- If the managed Microsoft server is a domain controller, open **Active Directory Users and Computers**, select the domain name, click **Builtin**, and add the user account to the Administrators Group.

#### Setting Specific Group Memberships and Permissions

If your security policy precludes adding user accounts to the Administrators group, you can add the user account to individual groups and grant only the required permissions. For guidelines and more information, see the following:

<http://support.microsoft.com/kb/325349>

<http://support.microsoft.com/kb/914392>

To add the user account of the Grid member to individual groups and grant specific permissions:

- To enable the member to synchronize DNS data with the Microsoft server, add its user account to the DnsAdmins Group.
- To enable the member to synchronize DHCP data with the Microsoft server, add its user account to the Dhcp Administrators Group.
- To enable the Grid member to monitor, start, and stop the DNS and DHCP services, grant the user account permissions on the Service Control Manager (SCM), as follows:
  - a. Grant permissions to the SCM on each managed Microsoft server. For more information, refer to the section *DNS Server Service Permissions* at <http://technet.microsoft.com/en-us/library/gg638675.aspx>.  
To find additional information, you can also search for "Least Privilege Setup" on the Microsoft sites.
  - b. Grant permissions to the DNS and/or DHCP service on each managed server by doing one of the following:
    - Use the `sc` command line utility to remotely configure each managed DNS or DHCP server. Note that you need to know the SID of the user account and its current permissions. You can retrieve the SID of the user account by using the `dsquery` and `dsget` commands.
    - Use the Domain Controller Policy editor to define a global policy that applies to all DNS or DHCP services running in a domain or on domain controllers. For additional information, refer to <http://support.microsoft.com/kb/324802>.

## Configuring a Managing Member

When you configure a member to manage Microsoft servers, you must specify the following:

- The management mode of the Microsoft server. For information, see [Setting the Management Mode](#) in the next section.
- A network view, if there is more than one in the Grid, and a DNS view, if there is more than one in the network view. For information, see [Synchronizing to a Network View and DNS View](#) below.

For the steps on configuring the managing member, see [Assigning Grid Members to Microsoft Servers](#) below.

## Setting the Management Mode

A Grid member can manage a Microsoft server in Read-only mode, which is the default, or in read-write mode. In Read-only mode, the Grid member copies the DNS and DHCP data from the Microsoft server to the Grid so Grid Manager admins can view the synchronized data. They cannot update the data, control the DNS and DHCP service of the Microsoft server, or configure any properties.

When you select Read-only mode for Active Directory sites, you can view the sites and networks that are present on the Microsoft server through Grid Manager. Note that you cannot manage the Active Directory sites and networks directly from the Grid, but you can manage an object within the Grid that is associated with a Read-only Active Directory Site or an Active Directory network. The synchronization process is Read-only and you cannot write into the Microsoft server in this mode. For more information, see [Assigning Grid Members to Microsoft Servers](#) below.

In Read/Write mode, Grid Manager admins are allowed to update the data of the Microsoft server. Therefore during each synchronization, the Grid member applies changes from the Grid to the Microsoft server and vice versa. Read/Write mode also allows admins to control DNS and DHCP services of the Microsoft server and configure some of their properties.

When you select Read/Write mode for Active Directory Sites, you can view and manage the sites and networks that are present on the Microsoft server through Grid Manager. When you update an object that is associated with the Active Directory Site or an Active Directory network, the changes reflect on the Microsoft server. For more information, see [Assigning Grid Members to Microsoft Servers](#) below.

Note that the management mode of a Microsoft server is separate from the admin permissions that the appliance requires to access the Microsoft servers and DNS and DHCP resources. An admin must still have the applicable permissions to the Microsoft servers and DNS and DHCP resources they want to access.



## Synchronizing to a Network View and DNS View

A Microsoft server can synchronize its data only to a single network view and a DNS view. Grid Manager automatically assigns Microsoft servers to the default view when a Grid contains only the default network view and DNS view. If a Grid has more than one network view, you must select a network view for the Microsoft server to synchronize its data; and if there are multiple DNS views, you must select a DNS view as well.

You cannot modify the assigned network view or DNS view of a Microsoft server after its data has been synchronized. Instead, you must remove the Microsoft server and then add it again. For information about removing a server, see [Removing a Managed Microsoft Server](#).

Microsoft servers do not support network views and DNS views. Therefore, network view and DNS view properties have no effect on the DNS and DHCP data that are synchronized from Microsoft servers.

## Assigning Grid Members to Microsoft Servers


To configure a Grid member to manage one or more Microsoft servers:

1. **Grid:** From the **Grid** tab -> **Microsoft Servers** tab -> **Servers** tab, click the Add icon.  
**Standalone appliance:** From the **System** tab -> **Microsoft Servers** tab -> **Servers** tab, click the Add icon.
2. In the *Add Microsoft Server(s)* wizard, complete the following:
  - **Which features do you want to configure?:** This section appears only when you have selected the **Enable MS AD feature** checkbox for mapping network users. For more information, see [Enabling Identity Mapping](#). You can select multiple options in this section:
    - **Network Users:** Select this checkbox to enable the Grid member to synchronize user information with the managed Microsoft servers.
    - **DNS and DHCP Services:** Select this checkbox to enable the Grid member to synchronize DNS and DHCP services with the Microsoft servers.
    - **Active Directory Sites:** Select this checkbox to enable the Grid member to synchronize Active Directory sites.
  - In the **General Settings** section, complete the following:
    - **Managing Member:** Click **Select Member** and select the Grid member that manages Microsoft servers. Select **None** if you do not want to associate a Microsoft server with a Grid member.
    - **Credentials to Connect to the Microsoft Server(s):** Enter the login name and password that the appliance uses to connect to the Microsoft servers. These must be the same as those you specified when you created the user account for the Grid member on the Microsoft servers. Note that you must specify the domain name and the user name in the following format: *domain\_name\user\_name*.
    - **Manage Server(s) in:** Select the management mode, which is either **Read-only** or **Read/Write**. You can choose to manage the DNS and DHCP synchronization services in either **Read-only** or **Read/Write** mode. For more information, see [Setting the Management Mode](#).
    - **Minimum Synchronization Interval (min):** The default synchronization interval is two minutes. This is the time between the completion of one synchronization and the start of a new one. Synchronizing large data sets could take longer than the synchronization interval, causing a delay in the start of the next synchronization. For example, if the synchronization interval is two minutes but a synchronization takes five minutes, the time between the start of the first synchronization and the start of the next one is approximately seven minutes. Note that the synchronization of Microsoft DHCP servers running Microsoft Windows 2012 or later includes the synchronization of DHCP failover relationships. Note that the DNS and DHCP failover synchronization rules do not have an impact on the Microsoft servers running a Windows version that is earlier than 2012.
  - **Logging Level:** Select a logging level for the Microsoft server log from the drop-down list: **Low**, **Normal**, **High**, and **Debug**. NIOS logs the messages based on the logging level you set.
    - **Low:** Logs only error messages.
    - **Normal:** Logs warning and error messages.
    - **High:** Logs warning, error and information messages.
    - **Debug:** Logs messages about all events associated with synchronization.



- See [Viewing Synchronization Logs](#) for a description of each level.
  - **Logging output destination:** From the drop-down list, select an output destination to which the appliance saves log messages for Microsoft servers. When you select **Microsoft Log**, the appliance logs the messages that are generated for the respective Microsoft server in the existing Microsoft log. This is selected by default. For more information, see [Viewing Synchronization Logs](#). When you select **Syslog**, NIOS logs the messages that are generated for the respective Microsoft server in the syslog. For more information about the syslog, see [Viewing the Syslog](#).
  - **Synchronize Data into Network View:** This field appears only when there is more than one network view in the Grid. Specify to which network view the data from the Microsoft servers is synchronized.
  - **Synchronize DNS Data into DNS View:** This field appears only when there is more than one DNS view in the network view. Specify to which DNS view the data from the Microsoft servers is synchronized.
  - **Comment:** You can enter additional information about the servers.
  - **Disable Synchronization:** Select this to disable the Microsoft servers. This allows you to preprovision the Microsoft servers and then enable them at a later time.
3. Click **Next**.  
Depending on your configuration in the **Which features do you want to configure?** section, the Add Microsoft Server(s) wizard displays the Microsoft server setting options.
4. Complete the following:
- If you have selected the **Network Users** checkbox, complete the following in the **Select your across-server settings for Network Users** page:
    - **Use General credentials (from first page of wizard):** Select this checkbox if you want to use the same credentials that you specified for connecting the Microsoft servers.
    - **Credentials for synchronizing Network User service information:** Specify a username and password to synchronize user information from Active Directory domain controllers. The username you specify here must belong to the Domain User group and Event Log Reader group in Microsoft. For information, see [Prerequisites on the Microsoft Server](#).
    - **Use General synchronization interval (from first page of wizard):** Select this checkbox to use the same synchronization interval that you specified in the **Minimum Synchronization Interval** for synchronizing the user and device mapping information from the Microsoft Active Directory authentication logs.
    - **Minimum synchronization interval:** The default synchronization interval is two minutes. This is the time between the completion of one synchronization and the start of a new one. Specify an interval to synchronize user information from the Microsoft Active Directory authentication logs.
  - If you have selected the **DNS and DHCP Services** checkbox, complete the following in the **Select your across-server settings for DNS and DHCP Services** page:
    - **Use General credentials (from first page of wizard):** Select this checkbox if you want to use the same credentials that you specified for connecting the Microsoft servers.
    - **Credentials to connect to DNS and DHCP Services:** Specify a username and password to synchronize DNS and DHCP services. You must use the same username and password that you specify here when the appliance prompts for credentials during DNS or DHCP synchronization.
    - **Use General synchronization interval (from first page of wizard):** Select this checkbox to use the same synchronization interval that you specified in the **Minimum Synchronization Interval** for synchronizing the DNS and DHCP services as well.
    - **Minimum Synchronization interval:** The default synchronization interval is two minutes. This is the time between the completion of one synchronization and the start of a new one. Specify an interval to synchronize the DNS and DHCP data from the Microsoft server.
    - **Manage DNS and DHCP services in:** Select a value from the drop-down list. You can choose to manage the DNS and DHCP synchronization services in either **Read-only** or **Read/Write** mode. For more information, see [Setting the Management Mode](#) above.
  - If you have selected the **Active Directory Sites** checkbox, complete the following in the **Select your across-server settings for Active Directory Sites** page:
    - **Use General credentials (from first page of wizard):** Select this checkbox if you want to use the same credentials that you specified for connecting the Microsoft servers. Clear the checkbox to specify a new username and password for managing Active Directory sites.
    - **Credentials for synchronizing Active Directory information:** Specify a username and password to synchronize Active Directory sites. You must specify the same username and password that you specify here when the appliance prompts for credentials while synchronizing Active Directory sites.

- **Use General synchronization interval (from first page of wizard):** Select this checkbox to use the same synchronization interval that you specified in the **Minimum Synchronization Interval** for synchronizing Active Directory sites.
  - **Minimum Synchronization interval:** The default synchronization interval is two minutes. This is the time between the completion of one synchronization and the start of a new one. Specify an interval to synchronize the Active Directory sites.
  - **Manage Active Directory sites in:** Select a value from the drop-down list. You can choose to manage the Active Directory Site in either **Read-only** or **Read/Write** mode. For more information, see [Setting the Management Mode](#) above.
  - **Encryption:** You can encrypt the network traffic between the Grid member and the managed Microsoft server using SSL. Select a value, **None** or **SSL**, from the drop-down list. Infoblox strongly recommends that you select **SSL** from the drop-down list to ensure the security of all communications between the NIOS appliance and the Active Directory server. When you select **SSL**, the appliance automatically updates the TCP port to 636. When you select this option, you must specify the FQDN of the Microsoft server instead of the IP address and you must upload a CA certificate from the Active Directory server. Click **CA Certificates** to upload the certificate. In the *CA Certificates* dialog box, click the Add icon, and then navigate to the certificate to upload it.
  - **TCP port for LDAP connections:** The appliance displays the port number by default based on the encryption type that you select. When you select **None**, the appliance automatically updates the TCP port to 389.
5. Click **Next** and do the following in the Managed Servers table:
- **Name or IP Address:** Enter either the FQDN or IP address of the Microsoft server. In order for the member to resolve the FQDN of a Microsoft server, you must define a DNS resolver for the Grid member in the **DNS Resolver** tab of the *Member Properties* editor. Note that if the IP address of the Microsoft server is specified, then the DNS resolver must resolve it when the member and Microsoft server synchronize DHCP data only.
  - **DNS Sync:** Select this option to enable the Grid member to manage the DNS service and synchronize DNS data with this server. Clearing this checkbox disables DNS service management and data synchronization. This allows you to pre-provision specific Microsoft servers and then enable them at a later time.
  - **DHCP Sync:** Select this option to manage the DHCP service of the Microsoft server and synchronize DHCP data with this server. Clearing this checkbox disables DHCP service management and data synchronization. This allows you to pre-provision specific Microsoft servers and then enable them at a later time.
  - **Active Directory Sites:** Select this option to manage Active Directory sites and synchronize Active Directory Sites and networks with the Grid.
  - **DNS Monitor & Control:** Click **Override** to override the setting inherited from the Grid. To inherit the same settings as the Grid, click **Inherit**. Select this to enable monitoring and the ability to control DNS service for the Microsoft server. For more information, see [Setting Grid Properties for Managing Microsoft Servers](#).
  - **Synchronize DNS Reporting Data:** Click **Override** to override the settings that are inherited from the Grid. To retain the same settings as the Grid, click **Inherit**. Select this to synchronize DNS reporting data from the Microsoft server. For more information, see [Synchronizing DNS Reporting Data](#). Note that synchronization of DNS reporting data is effective only when **DNS Sync** option is enabled for the Microsoft server.
  - **DHCP Monitor & Control:** Click **Override** to override the setting inherited from the Grid. To inherit the same settings as the Grid, click **Inherit**. Select this to monitor and control DHCP service for the Microsoft server. For more information, see [Setting Grid Properties for Managing Microsoft Servers](#).
- Note that you cannot start or stop a DNS or DHCP service on a specific Microsoft server if you disable the monitor and control setting for the respective service. You can control and monitor DNS and DHCP services at the Grid level and override the settings at the Microsoft server level. Each monitor and control setting applies only to the DNS or DHCP service and the respective Microsoft server.
- **Synchronize Network Users:** Click **Override** to override the settings inherited from the Grid. To inherit the same settings as the Grid, click **Inherit**. Select this to enable the identity mapping for the Microsoft server. For information, see [Enabling Identity Mapping](#). You can assign multiple Microsoft servers to a Grid member and test their connection to the Grid member. Click the Add icon to add another Microsoft server.

6. Select a Microsoft server and click the Test Microsoft Server icon, or click the Action icon  next to the respective Microsoft server and select **Test Microsoft Server** from the menu to verify whether the appliance can successfully connect to the Microsoft server. The appliance displays the test results in the *Test Microsoft Server Results* dialog box.
7. Save the configuration and click **Restart** if it appears at the top of the screen.

or

Click **Next**: Continue to the next step and define extensible attributes for the Microsoft servers. For information, see [Managing Extensible Attributes](#).

After you configure a Grid member to manage a Microsoft server, the member automatically connects to the Microsoft server and starts synchronizing data. You can then do the following:

- View the status of the servers in the *Microsoft Servers* panel, as described in [Monitoring Managed Microsoft Servers](#). Newly added servers first display a status of **Connecting** as the Grid member contacts the Microsoft servers. The status changes to **OK** after the Grid member successfully connects to the Microsoft server.
- View the data synchronized from the Microsoft servers. To view DNS data, navigate to the DNS view you specified. For information, see [Viewing Zones](#). To view DHCP data, navigate to the **Networks** tab of the network view that you specified. For information, see [Managing IPv4 DHCP Data](#).

Network conditions and the amount of data can affect the synchronization time. Therefore, you might not be able to view all of the synchronized data immediately.

- Use Smart Folders to organize the Microsoft servers and their data. For example, you can create a folder for DNS zones and another folder for DHCP scopes synchronized from a Microsoft server. For information about Smart Folders, see [Smart Folders](#).
- Update the synchronized data. For information, see [Managing Microsoft DNS Services](#), and [Managing Microsoft DHCP Services](#).

You can also use Global Search to search for synchronized data, such as zones and IP addresses. For information, see [Using Global Search](#).

## Configuring Grid Properties for Managing Microsoft Servers

You can configure the following Grid properties for Microsoft servers that are managed by a Grid member:

- Monitor and control settings for DNS and DHCP services. For information, see [Defining Monitor and Control Settings for DNS and DHCP Services](#) below.
- Use the identity mapping feature to get visibility of user interaction with their environments. For information, see [Configuring Identity Mapping](#).
- Output destination for Microsoft server log messages. For information, see [Output Destination for Log Messages](#) below.
- Simultaneous connections for Microsoft servers. For information, see [Maximum Simultaneous Connections for Microsoft Servers](#) below.
- Maximum connection timeout setting. For information, see [Configuring RPC Timeout Settings](#) below.
- Forward WINS packets to dedicated Microsoft Windows DNS and DHCP servers. For information see [Forwarding WINS Packets to Microsoft Servers](#) below.

For more information about configuring other Grid properties for Microsoft servers, see [Setting Grid Properties for Managing Microsoft Servers](#) below.

### Defining Monitor and Control Settings for DNS and DHCP Services

You can enable or disable monitor and control settings of DNS and DHCP services for a specific Microsoft server. The appliance enables this by default when you add a Microsoft server. When you upgrade the existing Microsoft servers, the managed member inherits values from the Grid. You can monitor and control the DNS and DHCP services on a Microsoft service only if both the management setting of the respective service and the monitor and control settings of the corresponding Microsoft server for the selected service are enabled. To know more about how to enable monitor and control settings, see [Setting Grid Properties for Managing Microsoft Servers](#) below.

 **Note**

The original setting that controls the overall management of a given service is referred to as the management setting. It controls whether the synchronization of the corresponding service is enabled or not, with no change to the existing synchronization behavior. Note that synchronization does not depend on the value of the monitor and control setting for the Microsoft server.

You can configure Microsoft server settings at the Grid level. Note that Microsoft servers inherit these settings by default, and you can override these settings at the Microsoft server level.

When you enable monitor and control settings for DNS and DHCP services, the managing member verifies the corresponding service status on the Microsoft server every 30 seconds. The Grid Master is notified of the status through Grid replication.

When you disable monitor and control setting for DNS and DHCP services, the managing member stops verifying the service status. NIOS administrators cannot start or stop DNS or DHCP service on the Microsoft server. When you try to start or stop these services through the Infoblox API, the appliance generates an error message. The pending service control requests made before disabling the monitor and control settings are sent to the Microsoft server.

For information about the displayed status, see [Viewing DNS and DHCP Service Status on Microsoft Servers](#).

### Synchronizing DNS Reporting Data

You can enable synchronization of DNS reporting data from the Microsoft server and view both NIOS and Microsoft services data in a single view in the supported DNS reports. This feature is enabled by default when you add a Microsoft server to the Grid. This feature is supported for Microsoft Windows 2012 R2 and Microsoft Windows 2016 versions, and you must disable this feature for all previous versions of Microsoft Windows servers. When you upgrade the Grid to NIOS 8.2 or later, this feature is disabled. For information about how to enable synchronization of DNS reporting data, see [Setting Grid Properties for Managing Microsoft Servers](#) below.

You can enable the synchronization of DNS reporting data at the Grid level. The Microsoft servers inherit these settings by default, and you can override these settings at the Microsoft server level. The synchronization of DNS reporting data is effective only if a valid Reporting license is installed on the Grid and when **DNS Sync** is enabled for the Microsoft server to synchronize the DNS data. For information about how to enable **DNS Sync**, see [Assigning Grid Members to Microsoft Servers](#).

 **Note**

You must enable DNS logging feature on the Microsoft server for this feature to function properly. To enable DNS logging feature on Microsoft servers, refer to <https://technet.microsoft.com/en-us/library/dn800669#some>.

When you enable the synchronization of DNS reporting data, the appliance performs the reporting data synchronization from the Microsoft server based on the specified time interval. The default synchronization interval is 15 seconds. You can change the synchronization interval using the CLI command `set ms_dns_reports_sync_interval`. For information, refer to the *Infoblox CLI Guide*. The collection of reporting data is dependent on the synchronization intervals set for the Microsoft server. Hence, there would be differences in the Microsoft services data updated in the reports as compared to the NIOS reporting data. Note that only events that are logged in the Microsoft event log are displayed in the Microsoft DNS reports. For a list of DNS reports that display data from both NIOS and the Microsoft server when this feature is enabled, see [Reports with Data Synchronized from Microsoft Servers](#).

Microsoft provides enhanced DNS logging and diagnostics for Microsoft Windows Server 2012 R2 and later versions. This includes DNS analytic events logging which enables activity tracking on the DNS server. An analytic event is logged each time the server sends or receives DNS information. In order to install enhanced DNS logging and diagnostics feature in Microsoft Windows Server 2012 R2 and later versions, you must apply the query logging and change auditing hotfix. Click the following link to apply the query logging and change auditing hotfix: <https://support.microsoft.com/en-us/help/2956577/update-adds-query-logging-and-change-auditing-to-windows-dns-servers>.

Note that DNS analytic events logging is not enabled by default on Microsoft servers. To install and enable DNS analytic events logging feature on Microsoft servers, refer to <https://technet.microsoft.com/en-us/library/dn800669#some>.

## Output Destination for Log Messages

You can configure an output destination for Microsoft server log messages at the Grid level and override it at the Microsoft server level. You can choose to save log messages related to Microsoft synchronization in the syslog or Microsoft log. For information, see [Viewing the Syslog](#).

When you change the setting, the Grid Master notifies the managing member about the new setting through Grid replication and sends log messages to the selected destination. The new setting takes effect for synchronization only after the managing member is notified. The synchronization, which is already in progress continues to log messages to the destination based on the old setting. To know more about how to select output destination for log messages, see [Setting Grid Properties for Managing Microsoft Servers](#) below.

## Maximum Simultaneous Connections for Microsoft Servers

You can specify a maximum number of simultaneous RPC connections that can be configured for the respective Microsoft server, which are managed by the Grid. You can override this value at the Microsoft server level. To know more about how to define maximum simultaneous connections for Microsoft servers, see [Setting Grid Properties for Managing Microsoft Servers](#) below.



### Note


When you increase the maximum number of simultaneous connections above the recommended setting for a given Microsoft server, it may consume additional bandwidth, memory, and CPU usage.

## Configuring RPC Timeout Settings

You can specify an RPC (Remote Procedure Call) timeout value to control network connectivity for Microsoft servers. You can configure an RPC timeout value at the Grid level and override it for each Microsoft server. The default value at the Grid level is ten seconds and the member inherits this value from the Grid. You can specify a value between one and 60 seconds. To know more about how to configure RPC timeout settings, see the next section, [Setting Grid Properties for Managing Microsoft Servers](#).

## Setting Grid Properties for Managing Microsoft Servers

To configure Grid properties for managing Microsoft servers, complete the following:

1. **Grid:** From the **Grid** tab -> **Grid Manager** tab, expand the Toolbar and click **Grid Properties** -> **Edit**. Select **Microsoft Sync Settings** tab in the *Grid Properties Editor* wizard.  
**Microsoft Server:** From the **Grid** tab -> **Microsoft Servers** tab -> **Servers** tab, select a Microsoft server and click the **Edit** icon, or click the Action icon  next to the respective Microsoft server and select **Edit** from the menu. In the Microsoft server editor, click the **General** tab.  
**Standalone appliance:** From the **System** tab -> **System Manager** tab, expand the Toolbar and click **System Properties** -> **Edit**.
2. Complete the following in the **Basic** tab:
  - **Logging output destination:** From the drop-down list, select an output destination to which the appliance saves log messages for Microsoft servers. When you select **Microsoft Log**, the appliance logs the messages that are generated for the respective Microsoft server in the existing Microsoft log. This is selected by default. For more information, see [Viewing Synchronization Logs](#). When you select **Syslog**, NIOS logs the messages that are generated for the respective Microsoft server in the syslog. For more information about the syslog, see [Viewing the Syslog](#). Click **Override** to select an output destination to save the log messages at the member level.
  - **Monitor DNS and DHCP Services:** You can enable monitoring and control services for DNS and DHCP services at the Grid level and override the settings for each service at the Microsoft server level. This is



enabled, by default. Each monitoring and control setting applies only to the corresponding service and is applicable to the respective Microsoft server only.

- **Monitor and control DNS Services:** Select this to enable monitoring and the ability to control DNS service for the Microsoft server.
  - **Synchronize DNS Reporting Data:** Select this to synchronize DNS reporting data from the Microsoft server. Clearing this checkbox disables DNS reporting data synchronization.
  - **Monitor and control DHCP Services:** Select this to enable monitoring and the ability to control a DHCP service for the Microsoft server.
3. Optionally, select the **Microsoft Server Settings** tab in the *Grid Properties Editor* wizard and complete the following in the **Advanced** tab or click the **Advanced** tab in the **General** tab in a Microsoft server editor:
    - **Maximum simultaneous connections:** Specify a maximum number of simultaneous RPC connections that can be configured for the respective Microsoft server, which are managed by the Grid. The default is five. You can specify a value between two and 40.  
You can click **Override** at the member level to specify a new value. The **Override** button changes to **Inherit**. Click **Inherit** to inherit the value from the Grid.
    - **RPC timeout:** Specify the RPC timeout value in seconds to control the network communication timeout. The default is ten seconds. You can specify a value between one and 60.  
You can click **Override** at the member level to specify a new value. The **Override** button changes to **Inherit**. Click **Inherit** to inherit the value from the Grid.
  4. Save the configuration.

## Forwarding WINS Packets to Microsoft Servers

If your Infoblox Grid includes legacy Microsoft DNS and DHCP servers, you can configure NIOS to forward WINS packets to dedicated Microsoft servers. Infoblox provides the `set wins_forwarding` and `show wins_forwarding` CLI commands you can use to perform this task. For detailed information about these commands, see [Using the NIOS CLI](#). When you enable port redundancy in NIOS, the LAN1 and LAN2 ports are grouped into one logical interface. They share one IP address and appear as one interface to the network. If a link to one of the ports fails or is disabled, the appliance fails over to the other port, avoiding a service disruption. When you enable port redundancy, WINS packet forwarding is not supported on the LAN2 interface. You must use the LAN1 interface.

### Note

Ensure that port 137 is not used for any services in your Grid; otherwise you will not be able to configure the appliance to forward WINS packets to Microsoft DNS and DHCP servers. Likewise, if you have enabled this feature, you will not be able to configure port 137 for any other services in your Grid.

## Managing Microsoft Servers


After you configure Grid members to manage Microsoft servers, you can set certain properties and manage the servers as follows:

- Set server properties, as described in the next section, [Setting Microsoft Server Properties](#).
- Change the managing member or the management mode, as described in [Changing the Managing Member or Management Mode](#) below.
- Back up the synchronized data, as described in [Backing Up Synchronized Data](#) below.
- Disable synchronization with a Microsoft server, as described in [Disabling Synchronization](#) below.
- Remove a Microsoft server, as described in [Removing a Managed Microsoft Server](#) below.

### Setting Microsoft Server Properties

You can modify any of the Microsoft server properties you previously configured, except for the network view and DNS view. You can also set certain properties, including the logging level, extensible attributes, and administrative permissions. Extensible attributes and permissions apply to the data only when they are managed from Grid Manager. Extensible attributes and permissions are not synchronized to the Microsoft server.

To set the properties of a Microsoft server:

1. From the **Grid** tab, select the **Microsoft Servers** tab -> **Servers** tab, select a Microsoft server and click the Edit icon, or click the Action icon  next to the respective Microsoft server and select **Edit** from the menu.
2. In the *Microsoft Server Properties* editor, you can set properties in the following tabs:
  - **General**: Modify the settings described in [Assigning Grid Members to Microsoft Servers](#).
  - **Services (DNS/DHCP)**: Modify DNS and DHCP synchronization settings. For more information, see [Assigning Grid Members to Microsoft Servers](#).
  - **Active Directory Domain/sites**: Modify Active Directory Site settings. For more information, see [Assigning Grid Members to Microsoft Servers](#).
  - **Extensible Attributes**: Define extensible attributes for the server. For information, see [Managing Extensible Attributes](#).
  - **Permissions**: Define administrative permissions that apply to the server. For information see [About Administrative Permissions](#).
3. Save the configuration and click **Restart** if it appears at the top of the screen.

You can edit the General and Logging properties of multiple Microsoft servers at the same time by selecting the Microsoft servers and clicking the Edit icon. When Grid Manager displays the *Microsoft Server Properties* editor, it displays the values that the Microsoft servers have in common. If a property has multiple values, it indicates this. You can then change any of the values and when you click **Save**, Grid Manager applies your changes to all the selected Microsoft servers.

### Changing the Managing Member or Management Mode

You can change the managing member and the management mode of a Microsoft server.

If you change the managing member, the previous member aborts any ongoing synchronization, and the newly assigned member resumes the synchronization process.

Note that if you switch the managing member or change the management mode of a Microsoft server from Read/Write to Read-only, the Grid member reverts any changes that were made from Grid Manager since the last synchronization. For example, an admin adds a network and DHCP range for a scope. If another admin changes the management mode of the Microsoft server to Read-only before the next synchronization, the Grid member deletes the network and DHCP range at the next synchronization.

To change the member or management mode:

1. From the **Grid** tab, select the **Microsoft Servers** tab -> **Servers** tab, select a Microsoft server and click the Edit icon, or click the Action icon next to the respective Microsoft server and select **Edit** from the menu.
2. In the *Microsoft Server Properties* editor, select the **General** tab and do any of the following:
  - **Managing Member**: Click **Select Member** and select another Grid member.
  - **Manage Server(s) in**: Select either **Read-only** or **Read/Write**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

### Backing Up Synchronized Data

When you back up the Grid, it includes all managed Microsoft data. If you restore a backup, the data is restored on the Grid only. It is not synchronized to the Microsoft servers. When the Grid member synchronizes the data after the restore operation, it overrides the data on the Grid with the data from the Microsoft servers. For information about backing up and restoring data, see [Backing Up and Restoring Configuration Files](#).

### Disabling Synchronization

When you set the disable option, the Grid member completes any on-going synchronization and does not start a new one. Setting this option only affects data synchronization and does not affect the operations of the Microsoft server. Synchronization resumes when the Microsoft server is re-enabled.

To disable a Microsoft server:

1. From the **Grid** tab, select the **Microsoft Servers** tab -> **Servers** tab, select a Microsoft server and click the Edit icon, or click the Action icon next to the respective Microsoft server and select **Edit** from the menu.
2. In the **General** tab, select the **Disable Synchronization** option.
3. Save the configuration and click **Restart** if it appears at the top of the screen.



## Removing a Managed Microsoft Server

When you remove a Microsoft server from the Grid, the managing member stops any on-going synchronization and does not start a new one. If the Microsoft server served DNS, the synchronized DNS data remains unchanged in the Grid. If the Microsoft server served DHCP, then Grid Manager deletes all the DHCP ranges, leases, and fixed addresses associated with the server. It also deletes networks that were assigned only to the Microsoft server. It does not delete a network if it was assigned to other Microsoft servers as well.

Removing a managed Microsoft server from the Grid does not affect the operations of the Microsoft server.

To remove a managed server:

1. From the **Grid** tab, select the **Microsoft Servers** tab -> **Servers** tab, select a Microsoft server and click the Delete icon, or click the Action icon next to the respective Microsoft server and select **Delete** from the menu.
2. When the *Delete Confirmation* dialog box appears, click **Yes**.

For information about how removing a Microsoft server affects the synchronized DNS and DHCP data, see [Disabling and Removing Microsoft DNS Servers](#) and [Disabling and Removing Microsoft DHCP Servers](#).

## Configuring Active Directory Sites and Services

An Active Directory Site represents physical or logical sites that are defined on a Microsoft server. Each Active Directory Site is associated with an Active Directory Domain. A Microsoft administrator can associate multiple sites and networks with an Active Directory Domain. You cannot associate the same network with multiple Active Directory Sites. For more information about adding an Active Directory Site, see [Monitoring Managed Microsoft Servers](#).

You can synchronize Active Directory Sites, along with the associated networks, from the Active Directory server. The networks that synchronize from the Active Directory server appears in NIOS as a network or a network container. The Active Directory Site to which the network belongs is displayed as an extensible attribute and associated with the respective network. For more information, see [Managing Extensible Attributes and Associated Networks](#). When you create a new Active Directory Site on NIOS, the appliance synchronizes new Active Directory Sites and associated networks with the Active Directory server.

The appliance uses LDAP to communicate with the Microsoft server. For more information, see as described in [Configuring Microsoft Server and LDAP Connections](#). You can specify a site link for a given Active Directory Site. For more information, see [Configuring Server Site Links](#).

You can assign relevant permissions for Active Directory Domains at the Grid level, network view level, or at the object level. For more information, see [Configuring and Managing Server Permissions](#) below.

For more information on supported versions, see [Supported Windows Versions](#).

## Best Practices for Configuring Active Directory Sites and Services

- A Grid member configured to synchronize Active Directory Sites and networks of a Microsoft server uses system resources, CPU, memory, network, etc. The amount of consumed resources depends on the number of Microsoft servers that are managed by the Grid member, the amount of data for synchronization and the synchronization frequency. Infoblox recommends that the managing Grid member should not serve other protocols.
- The Grid Master might also be affected by the initial synchronization. The first synchronization replicates all Active Directory Sites and networks from the managing Grid member to the Grid Master.
- The Grid member always initiates the connection to the Microsoft server. Infoblox recommends that you use an encrypted LDAP connection between the Grid member and the Microsoft server. The appliance displays a warning message when you use a non-encrypted connection.
- During a scheduled full upgrade, the synchronization of a Microsoft server is deactivated until the managing Grid member completes its upgrade.
- Microsoft Windows 2003 does not support IPv6 networks. Hence, the appliance prevents you from performing the following operations on the NIOS appliance, which might result in a synchronization failure on the Microsoft server:
  - Assigning an Active Directory Site to an IPv6 network if the site belongs to an Active Directory Domain that is managed by a Windows 2003 server.

- Creating an IPv6 network from an IPv6 network template in which an Active Directory Site belonging to an Active Directory Domain managed by a Windows 2003 server is defined.
- Moving one or more IPv6 networks when the destination Active Directory Site belongs to an Active Directory Domain that is managed by a Windows 2003 server.

## Synchronizing Active Directory Domains on a Domain Controller

The top level container in the Active Directory is called a Forest. Each forest can contain one or more Active Directory domains. All these domains share the same sites and networks. To avoid any inconsistency within NIOS, the appliance synchronizes only the root Active Directory domain of the forest into the Grid. For example, if the Microsoft server points to a domain controller of xyz.abc.com domain, the appliance synchronizes abc.com domain in NIOS.

An Active Directory user with the Domain User privilege (configured for synchronizing domains, sites, and their associated subnets) can read data from Active Directory and the same user with complete privileges on sites can write data from NIOS to Microsoft. However, these privileges are not enough to read deleted subnet associations of a site, when the delete operation occurs on the Microsoft server. For example, if you delete 10.x.x.0/24 from a Microsoft site and add 10.x.x.0/25 and 10.x.x.128/25 to the same site, after a synchronization, Grid Manager reflects all three subnets on the site.

NIOS performs incremental LDAP queries to search for the `uSNCchanged` attribute on objects. This fetches the newly added or modified objects but does not fetch deleted objects. To find these objects, a query must be run against the Deleted Objects container. Only users with the Domain Admin privilege can read from the Deleted Objects container using LDAP queries.

## Configuring and Managing Server Permissions

You can define permissions for Active Directory Sites on both the NIOS and Microsoft servers, but there is no relationship between these two sets of permissions and they are not synchronized. You must have relevant permissions on the Microsoft server to synchronize an Active Directory server. If you do not have permissions, the operation fails and the appliance logs the message into the Microsoft log. The synchronization process fails until appropriate permissions are granted.

On the Active Directory Sites, make sure that you set the following permissions in the *Permission Entry for Sites* dialog box:

- In the **Permissions** section, select all the checkboxes.
- In the **Properties** section, select the **Read All Properties** and **Write All Properties** checkboxes.

An administrator can define relevant permissions on the Active Directory Site and associated networks to prevent them from accidental deletion. However, you can modify such Active Directory Sites and networks. For example, you can rename an Active Directory Site.

You can define permissions for an Active Directory Domain at the following levels:

- Grid level: When you assign Grid level permission, it is applicable to the objects that are associated with the Active Directory Domains that you have defined. When you assign Grid permissions, you can assign either **Read/Write**, **Read-only**, or **Deny** permissions for **All Active Directory Domains** resource. For more information about how to define Grid level permission, see [Defining Global Permissions](#).
- Network view level: When you assign permission at the network view level, the permission is applicable to all the Active Directory Domains within the selected network view. For more information, see [Administrative Permissions for Network Views](#).
- Permission on the Active Directory Domain. For more information, see [Defining Object Permissions](#).

You must have the following permissions on the Active Directory Domain to perform the relevant operations:

- Read-only permission on the Active Directory Domain to view the Active Directory Domain and associated Sites.
- Read/Write permission on the Active Directory Domain to add, update or delete an associated Active Directory Site.
- Read/Write permission on the Active Directory Domain to associate or dissociate an Active Directory Site from or to the network.

- Note that the **All Active Directory Domains** permission will only support two modes: **Read-Only** and **Read/Write**. If you do not define permissions explicitly, the appliance sets Read-only permission for Active Directory Domains and Sites.
- Note that the extensible attributes are used to represent Active Directory Domains and Sites. As extensible attributes are generic and they do not support permissions, you can always retrieve Active Directory Domains and Sites for a given network if you have Read-only permission.

## Configuring Active Directory Sites and Associated Networks

An Active Directory Domain is a collection of Active Directory Sites and associated networks. Each Active Directory Site can have multiple networks associated with it. The synchronization process ignores a network which is not assigned to an Active Directory Site. You can create and modify Active Directory Sites and networks on the NIOS appliance.

You can use Active Directory Sites and services to synchronize sites and networks from the Active Directory server. After synchronization, the appliance displays these as networks or network containers in the appliance.

The Active Directory Sites that are associated with a network are displayed as extensible attributes, which is a combination of the Active Directory Domain and Site name, for the respective networks on the appliance.

To add Active Directory Sites, you must have Read/Write permission on the respective Active Directory Domain with which it is associated. For more information about permissions, see [Configuring and Managing Server Permissions](#).

To create Active Directory Sites and associate networks:

1. From the **Grid** tab, select the **Microsoft Servers** tab -> **Active Directory Domains** tab and click the respective Active Directory Domain name.
2. Click the Add icon in the *New Active Directory sites* wizard and specify a name for the Active Directory Site. You can associate multiple sites with an Active Directory Domain.



### Note

You can specify a name up to 63 bytes and it is not case-sensitive, but it cannot contain spaces and the following special characters: | { } ~ [ ] ' ; < > = ? @ ! " # \$ % ^ & ' ( ) + / \ , \* .

3. You can select an Active Directory Site and click the Delete icon to delete it.
3. Click **Next** to associate networks with an Active Directory Site. Select an Active Directory Site and click the Add icon to associate networks with the respective site.
4. Click **Cancel** to close the wizard without saving your settings. You can click **Save and Close** to save the settings and close the wizard or click **Save and Edit** to save the settings and edit the properties. The application will close the wizard and open the *Active Directory Site Properties* editor. Click **Save and New** to save the Active Directory Sites in the list and open a new wizard.

## Managing Active Directory Sites

You can manage Active Directory Sites and networks that you defined earlier and modify their information. You can do the following:

- View Active Directory Domains and sites, as described in [Viewing Active Directory Domains and Sites](#).
- Modify Active Directory Sites and networks, as described in [Monitoring Managed Microsoft Servers](#).
- Delete Active Directory Sites. For more information, see [Associating Active Directory Sites with Networks](#) below.
- Move networks that are associated with an Active Directory Site to another. For more information, see [Moving Networks Between Active Directory Sites](#) below.

### Viewing Active Directory Domains and Sites

You can view Active Directory Domains that are associated with the Microsoft server. You can also view Active Directory Sites and networks that are associated with the Active Directory Domains.

To view Active Directory Domains and sites:

1. From the **Grid** tab, select the **Microsoft Servers** tab -> **Active Directory Domains** tab.

2. Grid Manager displays the following information:
  - **Name:** The name of the Active Directory Domain. You can click the name to view the Active Directory Sites below it.
  - **NetBIOS Name:** The name returned in the NetBIOS format.
  - **MS Sync Server:** The Microsoft synchronization server that is associated with the Active Directory Domain.
  - **Network View:** The network view that is associated with the Active Directory Domain.You can do the following in the **Active Directory Domains** tab:
  - Sort the Active Directory Domains in ascending or descending order by column.
  - Use filters and the search function to look for specific values.
  - Export and print the information in the table.
3. To view Active Directory Sites associated with an Active Directory Domain, click the domain name that is displayed as a hyperlink. The appliance displays the list of Active Directory Sites that are associated with the respective Active Directory Domain in the **Active Directory Domains Home** table.

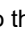
You can do the following in the **Active Directory Domains Home** table:

- Sort the Active Directory Sites in ascending or descending order by column.
- Use filters and the search function to look for specific values.
- Export and print the information in the table.

## Modifying Active Directory Sites and Networks

You can edit the name of an Active Directory Site and associate networks with the respective site. You can also add additional networks or remove associated networks from the Active Directory Site.

To edit an Active Directory Site:

1. From the **Grid** tab, select the **Microsoft Servers** tab -> **Active Directory Domains** tab and click respective Active Directory Domain name.
2. In the **Active Directory Domains Home** table, select an Active Directory Site name and click the Edit icon, or click the Action icon  next to the respective Active Directory Site name and select **Edit** from the menu.
3. In the *Active Directory Site Properties* editor, you can do the following:
  - **Name:** The name of the Active Directory Site. You can edit the name.
  - **Networks:** The networks that are associated with the Active Directory Site. You can click the Add icon to associate new networks with the Active Directory Site.  
To delete an associated network, select the checkbox adjacent to the network address, and click the Delete icon.
4. You can either click **Save and Close** or **Save** to save your settings. Click **Cancel** to close the editor without saving your settings.

## Deleting Active Directory Sites

You can delete Active Directory Sites that are associated with an Active Directory Domain. You cannot delete an Active Directory Site from the Active Directory Domain if the respective Active Directory Site has networks associated with it. You must first delete the associated networks to delete an Active Directory Site.

To delete Active Directory Sites:

1. From the **Grid** tab, select the **Microsoft Servers** tab -> **Active Directory Domains** tab and click the respective Active Directory Domain name.
2. In the **Active Directory Domains Home** table, select an Active Directory Site and click the Delete icon, or click the Action icon next to the respective Active Directory Site name and select **Delete** from the menu.
3. The appliance displays a confirmation message. Click **Yes** to delete the sites or click **No** to cancel the operation.

## Moving Networks Between Active Directory Sites

You can move networks from an Active Directory Site to another using either the **Microsoft Servers** or the **IPAM** tab. You can use the **Microsoft Servers** tab to move networks only when you select an Active Directory Site.

You can add multiple networks from various Active Directory Sites and move them simultaneously to an Active Directory Site. When the Active Directory Site synchronizes, it always maps to a certain network view. You can only select

networks from a network view to which the Active Directory Domain belongs.

To move networks within Active Directory Sites, complete the following:

1. **Microsoft Servers tab:** From the **Grid** tab, select the **Microsoft Servers** tab -> **Active Directory Domains** tab and click respective Active Directory Domain name. In the **Active Directory Domains Home** table, click the Action icon next to the respective Active Directory Site name and select **Move Networks** from the menu.  
OR  
**IPAM tab:** From the **Data Management** tab, select the **IPAM** tab, select networks that you want to move and click **Move Networks** from the Toolbar.
2. In the *Move Networks* wizard, complete the following:
  - **Destination Active Directory Site:** Click **Select Site** to select the destination Active Directory Site. The appliance displays the *Microsoft Sites Selector* dialog box listing all the Active Directory sites in the Grid. It also displays the smart folders, if any. You can use the filter and **Go to** functions to find a specific Active Directory site. To select an Active Directory site, click the site name.
  - **Networks:** You can select networks that you want to move to another Active Directory Site.
    - **Go to:** Specify an IP address of the network that you want to move to another destination and click **Go**.
    - Click the Add icon to add networks. To delete a network from the list, select the checkbox next to the network and click the Delete icon.
3. Click **Move** to move networks or click **Close** to exit.

### Moving Multiple Networks to an Active Directory Site

You can move multiple networks to a given Active Directory Site in a single operation. The following rules are applicable:

- Only a super-user can move multiple networks to an Active Directory Site.
- The destination site must belong to an Active Directory Domain that is synchronized in Read/Write mode.
- The appliance assigns the destination Active Directory Site to a network in either of the following cases:
  - If the network is currently assigned to another Active Directory Site within the same Active Directory Domain.
  - If the network is currently not assigned to any Active Directory Site within the Active Directory Domain.
  - The Active Directory Sites associated with an Active Directory Domain is different from the destination Active Directory Site.

### Associating Active Directory Sites with Networks

You can associate an Active Directory Site with an IPv4 or IPv6 network or an IPv4 and IPv6 network template. The appliance does not check the synchronization mode of the associated Active Directory Domain when you create or modify a template, but it verifies only when you use the template to create a network.

To assign Active Directory Sites to networks, complete the following:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab.
2. In the **Networks** section, select either **IPv4 Network** or **IPv6 Network** from the **Add** drop-down menu and complete the details as mentioned in [Adding IPv4 Networks](#) and [Adding IPv6 Networks](#) respectively, and complete the following to assign Active Directory Sites to networks:
  - **No Active Directory Site:** Select this if you do not want to associate an Active Directory Site with the network.
  - **Assign these Active Directory Domains/sites:** Select this to assign multiple Active Directory Sites to a network, but note that each Active Directory Site must be associated with a different domain. You cannot add an Active Directory Site, if you have already assigned a domain from the same Active Directory Site. When you select this option, the appliance enables the following:
    - **Active Directory Domain:** The Active Directory Domains that are synchronized from the Microsoft server. Click an Active Directory Domain that you want to associate.  
To search for a particular Active Directory Domain, specify the respective name and click **Go**.
    - **Active Directory Sites:** The Active Directory Sites that are associated with the selected Active Directory Domain. Click an Active Directory Site that you want to associate.  
To search for a particular Active Directory Site, specify the respective name and click **Go**.
    - **Add >:** Click this to add the selected Active Directory Sites to the network.

- **<Remove:** Click this to remove an Active Directory Site, which you have already added, from the network.
  - **<<Remove all:** Click this to remove Active Directory Sites, which you have already added, from the list.
  - **Assign the same domains/sites as the selected network:** Click **Select Network** to assign Active Directory Sites, which are associated with the selected network, to the new network that you are configuring. Select a network from the *Network Selector* dialog box. The appliance displays the Active Directory Domains and Sites that are associated with the selected network, but you cannot make any changes and the fields are greyed out. Click **Clear** to clear the entry.
3. You can either click **Save and Close** to save the settings and close the wizard or click **Save and New** to save your settings and open a new wizard or click **Cancel** to close the wizard without saving your settings.

## Managing Extensible Attributes and Associated Networks

The appliance creates an extensible attribute for each Active Directory Domain that is synchronized with the Grid. Note the following about extensible attributes:

- The appliance generates a name for the extensible attribute. The appliance displays an error message if you have already created an extensible attribute with the same name.
- Extensible attributes are Read-only and does not support inheritance.
- You can define a single value only and there are no default values.
- The appliance restricts an extensible attribute to IPv4 network and IPv6 network objects.
- You can specify a value for each extensible attribute and it is optional.

## Configuring Microsoft Server and LDAP Connections

Active Directory is a distributed directory service that authenticates network users. Active Directory uses LDAP (Lightweight Directory Access Protocol) to access other network services. This protocol is also used when an Active Directory Service entry is configured under Authentication Server Groups. For an Active Directory, the Grid member uses LDAP to communicate with the Microsoft server. Note that the managed Microsoft server must be a part of an Active Directory Domain and it must be a domain controller for the respective domain. The appliance logs an error message in the Microsoft log each time the synchronization happens if the Microsoft server is not a part of the Active Directory Domain or if it is not the domain controller for the respective domain. You can define an LDAP timeout value at the Grid level and override it at the Microsoft server level.

To configure LDAP timeout settings, complete the following:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties -> Edit** from the Toolbar. Select the **Microsoft Sync Settings** tab in the *Grid Properties Editor* wizard and click the **Advanced** tab.  
**Microsoft Server:** From the **Grid** tab -> **Microsoft Servers** tab -> **Servers** tab, select a **Microsoft server** and click the Edit icon. In the Microsoft server editor, select the **General** tab and click the **Advanced** tab.
2. Complete the following:
  - **LDAP timeout:** Specify the LDAP connection timeout value. The default is 10 seconds. You can specify a value between one and 60 seconds.  
You can click **Override** to override the LDAP timeout value at the Microsoft server level. To inherit the settings from the Grid, click **Inherit**.
3. Save the configuration.

## Configuring Server Site Links

You must choose a site link when you create an Active Directory Site on the Microsoft server. The site link describes how the replication occurs between the sites and the protocol, either IP (Internet Protocol) or SMTP (Simple Mail Transfer Protocol), which is used for communication.

The appliance uses the default object DEFAULTIPSITELINK when you create an Active Directory Site from the NIOS appliance. If you want to change the site link for a given Active Directory Site, which is already created, you must create the Active Directory Site on the Microsoft server or modify it on the Microsoft server after you create it using the NIOS appliance.

You must first configure site links on the Microsoft server. When you create Active Directory sites, the appliance automatically associates them with the default IP site link. The value that you specify for default IP site link must exist on the Microsoft server. Note that the appliance does not display any error message when you create a site with a site link

that does not exist on the Microsoft server.

To configure server site links, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.  
**Microsoft Server:** From the **Grid** tab -> **Microsoft Servers** tab -> **Servers** tab, select a **Microsoft server** and click the Edit icon. In the Microsoft server editor, select the **General** tab and click the **Advanced** tab.
2. Select the **Microsoft Sync Settings** tab in the *Grid Properties Editor* wizard and complete the following in the **Advanced** tab:
  - **Default IP site link:** Specify the default IP site link in the form of a string. The appliance does not validate it against the Windows server during configuration. The appliance displays an error message during synchronization, if the site link for IP does not match the configured name on the Windows server. You can click **Override** to override the default IP site link value at the Microsoft server level. To inherit the settings from the Grid, click **Inherit**.
3. Save the configuration.

## Managing Active Directory Sites and Associated Networks

You can manage Active Directory Sites and networks using either the Microsoft servers or the NIOS appliance.

### Managing an Active Directory Site from the Microsoft server

You can create, modify, or delete an Active Directory Site from the Microsoft server in both Read-only and Read/Write synchronization modes. When you create an Active Directory Site from the Microsoft server, it is created on the NIOS appliance. This site appears as an Active Directory Site under the corresponding Active Directory Domain and in the respective extensible attribute.

When you make changes on the Microsoft server, the synchronization process ignores any other modifications that you make, but replicates only the name changes to the Grid. When you rename an Active Directory Site, the change reflects:

- in the Active Directory Site and in the respective extensible attribute.
- on all the networks that are associated with the Active Directory Site.

When you delete an Active Directory Site from the Microsoft server, it deletes the Active Directory Site object and removes corresponding extensible attributes. The appliance removes the value from all the network objects that are associated with it.

### Managing an Active Directory Site from NIOS

You can create, modify, or delete an Active Directory Site from the appliance in Read/Write synchronization mode only. The appliance displays an error message if you perform any of these operations in Read-only mode.

When you rename an Active Directory Site and the associated networks, the changes reflect on the Microsoft server. The appliance displays an error message when you delete an Active Directory Site that is associated with networks. You must first remove the networks that are associated with the site.

You cannot resize or split a network with an associated site in Read-only synchronization mode. You cannot join two networks if any of the involved networks or network containers are associated with a site irrespective of the synchronization mode.

### Managing Network Operations from Microsoft server

Note that you can perform network operations, such as create, modify or edit networks, in both Read-only and Read/Write synchronization modes. When you create a network, the synchronization process ignores the network if it is not assigned to a Active Directory Site. When the network is attached to the site, the appliance finds a matching network or network container in the configured network view. If it does not exist, the appliance creates a new network or a network container depending on the objects existing in the database. The new network or network container is set as unmanaged and a new extensible attribute is attached to the object to store the site.

Note the following about managing networks:

- When you update a site that is associated with a network, the change is visible on the corresponding network or network container.



- Synchronization process ignores the changes when you update a network. Note that it does not ignore the changes when you modify the Active Directory Site or change an IP address or a CIDR.
- When you delete a network, the appliance removes the extensible attribute that is representing the associated site from the corresponding network or network container.
- If the status is set as managed, it indicates that the network is used for another purpose inside the Grid. The appliance does not remove it from the Grid.
- If the status is set as unmanaged, the appliance deletes the network if there is no other site coming from another Active Directory Domain, which is associated with the network. Otherwise the appliance does not remove the network from the Grid.

## Converting Unmanaged Networks to Managed Status

When you synchronize a network from the Active Directory server, the network created in NIOS is considered to be unmanaged. You can neither associate a NIOS member or a Microsoft server to manage it nor edit permissions, and other fields. You can only edit extensible attributes and Active Directory sites on the unmanaged network. To perform certain operations on Unmanaged networks, you must first convert them to Managed networks.

The IPAM main page lists all networks that are created during Active Directory synchronization as unmanaged, highlighted in yellow. You can explore unmanaged networks through IPAM's **IP Map** and **IP List** views.

Unmanaged networks can be converted at the IPAM main page and at the device level under **Data Management** → **Devices**, selecting a device and opening the **Networks** page.

## Managing Network Operations from NIOS

The appliance denies the operation when you create a network, which is associated with a site, if the respective site belongs to an Active Directory Domain that is synchronized in Read-only mode. Otherwise the appliance creates it on the Microsoft server that is associated with the respective site.

## Restoring Server Data

When you restore a database backup, the appliance performs the following operations after the standard restore procedure is complete:

- The appliance deletes the cached data that is used for synchronization.
- Synchronization for all configured Microsoft servers starts in Read-only mode. Note that these rules are applicable during upgrade also.

## Configuring Identity Mapping

You can enable Identity Mapping on the NIOS appliance to provide Active Directory domain user information if the NIOS appliance is connected to a Microsoft server. This feature supports Active Directory domains whose domain controller is running the supported Windows server. For more information on supported windows versions, see [Supported Windows Versions](#).

Note that Identity Mapping is not supported for the Windows 2003 server and earlier editions. Note that Identity Mapping is not supported on the IB-VM-810 and IB-VM-820 appliances.

Each network user being mapped can use different devices to access the Windows environment. So using the identity mapping feature and synchronizing all Microsoft servers on the Infoblox appliance provide visibility of user interaction with their environments. By enabling this feature, you can monitor Active Directory domain users, the IP addresses they log on to, the login status, and the time duration of their current status in the **IPAM** tab.

To view user information, you must first enable this feature at the Grid level. You can enable this feature even when you have not installed an **MS Management** license on the appliance. However, you cannot configure DNS, DHCP, and Active Directory sites synchronization unless you install an **MS Management** license on the appliance.

When you enable this feature, the appliance remotely communicates with all synchronized Microsoft servers (Domain Controllers, an Exchange server, or a domain member) to pull event logs. The identity mapping information displayed is as accurate as these event logs are available in the Microsoft authentication logs. Therefore, it is necessary to assign Grid members to Microsoft servers to collect user information from Windows event logs. For information, see [Assigning Grid Members to Microsoft Servers](#).

 **Note**

The identity mapping information displayed on NIOS completely depends on live event logs that are available on the Microsoft servers. The appliance pulls event logs incrementally. So subsequent requests pull only the latest logs since the last successful synchronization. To avoid data loss, depending on the expected activities, you must consider the size of the event log file on the Microsoft server and how often you want to synchronize the data with the appliance before the event log file rolls over.

## Prerequisites on the Microsoft Server

You must enable event logs on the Microsoft server for the Identity Mapping feature to function properly. To enable event logging on Microsoft servers, refer to <https://technet.microsoft.com/en-us/library/dd941595%28v=ws.10%29.aspx>.

The identity mapping information is collected successfully only when the Microsoft users belonging to a Domain User group and Event Log Reader group start a RPC session and access MS-EVEN6. The synchronization process is successful when they have this permission. The synchronization process fails unless appropriate permissions are granted. The failed operations are logged in the Microsoft logs. The NIOS appliance tries to collect user information again in the next synchronization.

## Administrative Permissions

Only superusers can view identity mapping information. Limited-access admin groups can view identity mapping information only if they have network permissions. For example, if the users have permissions to only DNS zones, they may not be able to view identity mapping information because they do not have network permissions. The appliance does not display a warning message if admins do not have correct permissions. For information about network permissions, see [Administrative Permissions for IPv4 and IPv6 Networks and Shared Networks](#).

## About User Sessions

A user session is an abstract concept to specify a single user logging with a network address for a finite period of time. A user session starts when a Kerberos account authentication event or login event is received and ends when a logout is received, although such an event may never be received. In such cases, a session is considered to be timed out. Network user associations are unique for a finite period of time. A single login involves a number of login and logout events. In order to consolidate and improve system performance, Infoblox uses the concept of consolidation window. If a login event or Kerberos service ticket is received within the consolidation window, then that session is considered as an extension from the previous session. If a login event is received outside this window, it is considered as a new session. The number of network users you see depends on the length of the consolidation window. For example, consider the following sample events occurred on the Microsoft server when the consolidation window is set to 10 minutes:

```
Kerberos Authentication Request: 10:00:00
```

```
Kerberos Service Ticket Request: 10:00:01
```

```
Account Logon: 10:00:02
```

```
Account Logoff: 10:00:03
```

```
Account Logon: 10:00:05
```

```
Account Logoff: 10:00:07
```

```
Account Logon: 10:00:10
```

```
Account Logoff: 10:00:11
```

The sample user session for the above events:

User Name	IP Address	First Seen Time (UTC)	Last Seen Time (UTC)
John	10.10.10.10	10:00:00 AM (UTC)	10:00:11 AM (UTC)

If another login request is received at 10:30 AM (10 minutes after the last seen event), then it is considered as a different session:

User Name	IP Address	First Seen Time (UTC)	Last Seen Time (UTC)
John	10.10.10.10	10:00:00 AM (UTC)	10:00:11 AM (UTC)
John	10.10.10.10	10:30:00 AM (UTC)	

If a Kerberos service ticket is received instead of a login event, then the previous session is extended and is updated as **Last Seen Time** in the first user session.

User Name	IP Address	First Seen Time (UTC)	Last Seen Time (UTC)
John	10.10.10.10	10:00:00 AM (UTC)	10:30:00 AM (UTC)

The appliance displays separate entries (counts) for the following scenarios:

- Multiple users logging in from the same IP address. For example, User 1 logged in with 10.10.10.10 address counts as one network user and User 2 logged in with 10.10.10.10 address counts as another network user.
- Same user logging in from multiple IP addresses. For example, a single user can log in to multiple workstations, each with a different IP address.
- Same user logging from the same IP address at different time intervals.

#### Network User Count Displayed for Different Login Scenarios

If you are configuring multiple Domain Controllers belonging to different organizations, then you must configure them in different network views. For example, if the same user logs in to multiple Domain Controllers with the same IP, it creates multiple entries for each login. In this case, there is a chance of overwriting an entry by subsequent events. If Domain Controllers are configured in different views, then separate entries are displayed for different network views. You can configure multiple Domain Controllers within the same network view if they serve the same organization, possibly using a load balancing method.



#### Note

The appliance displays user information only for the managed networks.

The following table illustrates how the appliance displays counts for different login scenarios:

#### *Network User Count Displayed for Different Login Scenarios*

Login Scenarios	Appliance Displaying User Mapping Information
Mobile user logging in to the Microsoft Exchange server	Note that you must first synchronize both the Domain Controller and Microsoft Exchange server with the appliance to get user mapping information for this scenario. For this example, two entries are displayed. <ol style="list-style-type: none"> <li>1. User name and IP address of the Microsoft Exchange server on the Domain Controller.</li> <li>2. User name and IP address of the mobile device on the Microsoft Exchange server.</li> </ol>
Multiple users from the same IP address	Appliance displays separate entry (user name and IP address) for each user.
Same user from multiple IP address	Appliance displays separate entry (user name and IP address) for each user.

## Login and Logout Timestamps

Note that all timestamps are displayed in the time zone of the admin account that you use to log in to the appliance. There is a possibility of missing the login and logout events as described in the following cases:

- There is a chance of missing a login event when NIOS retrieves event logs from the Microsoft server after a user logs in and the event log has already rolled over.
- There is a possibility of missing a logout or session end notifications when the user shuts down the workstation or leaves the system. In such cases, the Microsoft event log does not specifically indicate a logout. The logout events that are missed on the Domain Controller or Exchange server are missed on NIOS as well.

To maintain accuracy, the login timestamp is estimated as logout timestamp minus (-) the idle timeout. However, when a login or Kerberos Authentication event is received, the login timestamp is updated to the value available in the Kerberos authentication event or the login event.

To maintain accuracy of the logout time data, the appliance allows you to configure the length of idle time in the *Grid Properties Editor* wizard. After this time interval, the status of the user changes to **Timed Out**. For information about how to set timeout length, see [Configuring Active User Timeout Session](#).



### Note

The **Timed Out** and **Logged Out** user information is periodically removed from the database.

## Viewing Active Directory User Information

To view Active Directory user information, you must first enable identity mapping feature at the Grid level. For information about enabling Identity Mapping feature, see [Enabling Identity Mapping](#). After you enable the identity mapping feature, you must synchronize the appliance with all Microsoft servers in order for the appliance to gather user and device mapping information from the Microsoft servers. You can view Active Directory user information in the **Network Users** tab. For more information, see [Viewing Active Network Users](#).

To synchronize the appliance with Microsoft servers:

1. From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.
2. Select the **Microsoft Integration** tab in the *Grid Properties Editor* wizard, and complete the following:
  - **Synchronize Network Users with all MS servers:** Select this checkbox to synchronize users with all Microsoft servers that are managed by the Grid in order for the appliance to gather user and device mapping information from the Microsoft server authentication logs. You can override this value at the Microsoft server level.



### Note

On an Infoblox appliance, the **Enable Network Users Feature** and **Synchronize Network Users with all MS servers** options are disabled by default for all new installations.

3. Save the configuration.

## Monitoring Managed Microsoft Servers

You can monitor the status of managed Microsoft servers from the Dashboard and from various panels in the **Grid** tab. Grid Manager also maintains a log for each managed Microsoft server. You can monitor Microsoft servers and their services as follows:

- You can view the *Microsoft Servers Status* widget on the Dashboard. For information, see [Microsoft Servers Status Widget](#).
- You can view the status of Microsoft servers. For information, see [Viewing DNS and DHCP Service Status on Microsoft Servers](#) below.
- You can view the logs of the Microsoft servers. For information, see [Viewing Synchronization Logs](#) below.

### Viewing DNS and DHCP Service Status on Microsoft Servers

Each Microsoft server reports the following statuses:





- **Service status (DNS and DHCP):** Displays the status of the service on the actual Microsoft server, based on the latest polling of the service by the managing member. The monitor and control setting has an effect only on the service status, and therefore can affect the overall status.
- **Overall status:** Displays the service status for each service that is enabled for synchronization with the Microsoft server. The overall status ignores any service status for which the monitor and control setting is disabled.
- **Synchronization status:** Displays the synchronization status for each service that is enabled for synchronization on the respective Microsoft server. The synchronization status is not affected by the monitor and control setting.

#### Note





When you disable monitor and control settings, Grid Manager displays **unknown** using gray color for such services. When you enable monitor and control setting of a given service, Grid Manager displays the last known status that is obtained before the setting was first disabled. The appliance later updates to the latest status as soon as the monitoring resumes and Grid Manager displays the new status.

You can view details about the managed Microsoft servers by navigating to the **Grid** tab -> **Microsoft Servers** tab -> **Servers** tab. For each Microsoft server, the panel displays the following by default:

- **Name:** The FQDN of the Microsoft server.
- **Status:** The connection status, which can be one of the following:
  - **Running:** The Grid member is connected to the Microsoft server.
  - **Connecting:** The Grid member is connecting to the Microsoft server.
  - **Error:** The Grid member failed to connect to the Microsoft server. Check the Microsoft log for any messages to determine the reason for the failure.
  - **Unknown:** The Microsoft server is disabled. The Grid member does not try to connect to disabled servers.
- **IP Address:** The IP address of the Microsoft server.
- **DNS:** The status of the DNS service on the Microsoft server. When you disable DNS synchronization, NIOS does not display any status icon. The status icon can be one of the following:

Icon	Color	Meaning
	Green	The DNS service is functioning properly.
	Red	The Microsoft server is unavailable.
	Yellow	The DNS service is starting or stopping.
	Gray	The DNS service is stopped or management of the Microsoft DNS server is disabled.


- **DHCP:** The status of the DHCP service on the Microsoft server. When you disable DHCP synchronization, NIOS does not display any status icon. The status icon can be one of the following:

Icon	Color	Meaning
	Green	The DHCP service is functioning properly.
	Red	The Microsoft server is unavailable.
	Yellow	The DHCP service is starting or stopping.
	Gray	The DHCP service is stopped or management of the Microsoft DHCP server is disabled.

- **Comment:** Displays any comments that were entered for the Microsoft server.
- **Site:** Displays any values that were entered for this pre-defined attribute. You can add the following columns for display:
- **Version:** The Windows version of the managed server.
- **Managing Member:** The hostname of the Grid member that manages the server.
- **Synchronization Status:** Displays the synchronization status as follows:
  - **Running:** The Microsoft server is synchronizing data with the Grid member.
  - **Connecting:** The Grid member is trying to connect to the server.
  - **Error:** Synchronization failed between the member and server. You can check the messages in the Microsoft server log to determine the reason for the failure.

- **Last Changed:** Displays information about when the status was last updated for Microsoft DNS and DHCP services. It corresponds to the last time information was exchanged with the server.
- **AD Domain:** Displays the AD domain of the Windows server. This is displayed only if the Windows server belongs to an Active Directory domain.
- **Root AD Domain:** Displays the root AD domain of the Windows server. This is displayed only if the Windows server belongs to an Active Directory domain.

You can also do the following:

- Add Microsoft servers.
  - Click the Add icon.
- Edit the properties of a Microsoft server.
  - Click the checkbox beside a server and click the Edit icon, or click the Action icon  next to the respective Microsoft server and select **Edit** from the menu. For information, see [Setting Microsoft DHCP Server Properties](#).
- Delete a Microsoft server.
  - Click the checkbox beside a server and click the Delete icon, or click the Action icon next to the respective Microsoft server and select **Delete** from the menu. For information, see [Removing a Managed Microsoft Server](#).
- Manage DNS and DHCP services of a Microsoft server.
  - Click the checkbox beside a server and click the Manage Server Services icon, or click the Action icon next to the respective Microsoft server and select **Manage Server Services** from the menu to view the service status. You can mouse over the DNS and DHCP service icons and click the Start/Stop service icon to start or stop a service, or click the Edit Service icon to edit the service properties. For information about setting DHCP server properties, see [Setting Microsoft DHCP Server Properties](#). For information about setting DNS server properties, see [Specifying Forwarders for Microsoft Servers](#).
- View detailed server status information, as described in the next section Viewing Detailed Status Information.
- Click **Test MS Server** to test the Microsoft server connection. The appliance validates the Microsoft server and displays the test code and the test result data for services that you have enabled. To test the Microsoft server, click the checkbox beside a server and click the Test Microsoft Server icon, or click the Action icon next to the respective Microsoft server and select **Test Microsoft Server** from the menu. For more information, see [Assigning Grid Members to Microsoft Servers](#).
- View extensible attributes associated with the Microsoft server.
  - Click the Action icon next to the respective Microsoft server and select **Extensible Attributes** from the menu. For information, see [Assigning Grid Members to Microsoft Servers](#).
- Define permissions for Microsoft servers.
  - Click the Action icon next to the respective Microsoft server and select **Permissions** from the menu.
- Use filters and the **Go To** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Sort the displayed data in ascending or descending order by column.
- Export the list of Microsoft servers to a .csv file.
  - Click the Export icon.
- Print the list of Microsoft servers.
  - Click the Print icon.

## Viewing Detailed Status Information

You can view more status information by selecting a server from the Microsoft servers panel and clicking the Action icon next to the respective Microsoft server and selecting **Detailed Status** from the menu or clicking the Detailed Status icon. The **Detailed Status** panel displays the following information:

- **Synchronization Status:** The status icon indicates the synchronization status as follows:
  - Green: The Microsoft server is synchronizing data with the Grid member.
  - Red: Synchronization failed between the member and server. You can check the messages in the Microsoft server log to determine the reason for the failure.
- **DNS Service Status:** For information about the status icons, see [Viewing DNS and DHCP Service Status on Microsoft Servers](#) above.



- **DNS Service Status Last Updated:** The date and time of the last DNS service status update from the Microsoft DNS server.
- **DHCP Service Status:** For information about the status icons, see [Viewing DNS and DHCP Service Status on Microsoft Servers](#) above.
- **DHCP Service Status Last Updated:** The date and time of the last DHCP service status update from the Microsoft DHCP server.
- **Active Directory Sync Status:** The Active Directory Site is synchronizing data with the Grid member when the status icon is green.
- **Active Directory sync status last updated:** The date and time of the last Active Directory Site update from the Microsoft server.

Note the following guidelines about status information:

- Grid Manager does not display any status information if there is no synchronization between DHCP and DNS.
- If the appliance has not received any information when the services are enabled, then the Synchronization Status icon is displayed in red, whereas the DNS and DHCP status icons are displayed in grey.

## Viewing Synchronization Logs

Grid Manager maintains a synchronization log file for each Microsoft server managed by a Grid member. It logs events related to the synchronization process, depending on the logging level that you configured in the **Logging** tab of the *Microsoft Server Properties* editor described in [Setting Microsoft Server Properties](#).

The log files are rotated and compressed once they reach 40MB. To view the log file of managed Microsoft server:

1. From the **Administration** tab, select the **Logs** tab -> **Microsoft Logs** tab.
2. If there is more than one managed server in the Grid, you can select the Microsoft server whose logs you want to view.
3. The log file contains information related to the synchronization of the Microsoft DNS and DHCP data, as follows:
  - **Timestamp:** The date and time of the log message. The time zone is the time zone configured in the User Profile.
  - **Source:** Identifies the event that generated the message, such as a server synchronization or zone synchronization.
  - **Level:** Indicates the severity of the message, which can be one of the following:
    - **Debug:** Provides information about all events associated with synchronization.
    - **Information:** The Grid member is synchronizing with the Microsoft server and these messages provide normal status information.
    - **Warning:** The Grid member synchronized the data, but there was an issue, which is detailed in the Message section.  
If the Grid member encounters an error during the synchronization, it skips the object with the error, logs the error in the Microsoft log, and continues to synchronize the rest of the data. The Grid member logs the error at each synchronization until you resolve the issue and it can synchronize the object successfully.
    - **Error:** The Grid member failed to synchronize an object, such as a DNS zone or DHCP scope, due to the error described in the Message section.
  - **Object Type:** The type of object that corresponds to the entry, such as FQDN or ADDRESS.
  - **Object Name:** The name of the object that corresponds to the entry.
  - **Message:** Detailed information about the event.

You can also do the following in the log viewer:

- Toggle between the single line view and the multi-line view.
- Navigate to the next or last page of the file using the paging buttons.
- Refresh the view.
- Click the Follow icon to have the appliance automatically refresh the log every five seconds.
- Download the log.
- Clear the contents of the log.
- Sort the data in ascending or descending order by column.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Export or print the content of the log.

## Managing Microsoft DNS Services

After you configure a Grid member to manage a Microsoft DNS server, the Grid member connects to the Microsoft server and starts synchronizing DNS data from the Microsoft server to its database. First, it synchronizes the Microsoft server properties and its list of zones. Then it synchronizes each zone individually, including its properties and resource records.

The synchronization time varies, depending on different factors, such as the number of managed Microsoft servers and the amount of data being synchronized. The synchronized data is then replicated to the Grid Master through the Grid replication process.

If the server is managed in read/write mode, admins can update the synchronized DNS data, control the DNS service of the server, and specify forwarders for it as well.

This section provides guidelines for using Grid Manager to manage Microsoft DNS servers and for synchronizing DNS data between Microsoft servers and the Grid. It discusses some features of the Microsoft DNS servers only as they relate to the synchronization of data. Please review the Microsoft documentation for complete information about Microsoft DNS servers and their features.

In addition, if you encounter technical issues with your Microsoft DNS servers, contact Microsoft Technical Support or consult the Microsoft Support site at <http://support.microsoft.com/>. Some Windows versions require certain updates and hotfixes installed, so the Microsoft server can synchronize with the Grid member. For information about these requirements, see [Supported Windows Versions](#).

The topics in this section include:

- [Synchronizing DNS Data](#)
- [Enabling and Disabling DNS Zone Synchronization](#)
- [IDN Support for Synchronized DNS Data](#)
- [Managing Synchronized DNS Data](#)
- [Synchronizing Updates](#)
- [Viewing Members and Managed Servers](#)
- [Specifying Forwarders for Microsoft Servers](#)
- [Disabling and Removing Microsoft DNS Servers](#)

## Synchronizing DNS Data

Managing members synchronize the properties and resource records for the following types of DNS zones:

- Authoritative forward-mapping zones
- IPv4 and IPv6 reverse-mapping zones
- Stub zones
- Delegations
- Active Directory-integrated zones

Grid members synchronize most of the resource records supported by Microsoft servers, except for WINS, WINSR, and ATMA records. They synchronize all the resource records supported by Infoblox DNS servers, as well as unsupported records, such as ISDN and X25 records. You can view the unsupported records in Grid Manager and delete them, but you cannot edit them. Note that Grid Manager and Microsoft DNS servers display some resource records, such as SIG records, in a different format. You can enable and disable zone synchronization for individual Microsoft DNS zones. For information, see [Enabling and Disabling DNS Zone Synchronization](#).

Grid members do not synchronize the following DNS zones supported by Microsoft servers:

Infoblox Terminology	Microsoft Terminology
Forwarding zones	Conditional forwarders
Cached zones	Stub zones

Infoblox Terminology	Microsoft Terminology
Root zone	Root zone (Dot zone)
0.in-addr.arpa	0.in-addr.arpa (0.0.0.0)
127.in-addr.arpa	127.in-addr.arpa (127.0.0.1 - loopback)
255.in-addr.arpa	255.in-addr.arpa (255.255.255.255 - broadcast)
TrustAnchors	Trust Anchors

You cannot use Grid Manager to create the unsupported zones and assign them to a Microsoft server. Any zone on the Grid that has the same name as a forwarding, cached or root zone on the Microsoft server is not synchronized. In addition, Grid members do not synchronize the contents of a zone if the Microsoft server is a secondary server. Subdomains defined within a Microsoft DNS zone are not synchronized unless they contain at least one resource record. For example, in the corpxyz.com zone, any resource record defined in a subdomain of the corpxyz.com zone is synchronized. If the subdomain sub.corpxyz.com zone has no resource record, it is not synchronized. The following zones and resource records are supported on Microsoft servers running Windows Server 2008 only. Therefore, Grid members can only synchronize these DNS zones and resource records with Microsoft servers running Windows Server 2008.

- IPv6 reverse-mapping zones
- Global Names zones
- DNAME records
- NAPTR records
- DNSSEC records

## Synchronizing with Multiple Servers

Because a Grid member can manage multiple Microsoft servers, it could potentially manage multiple servers assigned to the same zone. For example, a Grid member could manage a Microsoft server that is the primary server of a zone and one or more Microsoft servers that are secondary servers of the same zone. It could also manage multiple Microsoft servers that are secondary servers for the same zone.

If a Grid member manages the primary server and at least one secondary server of a zone, the Grid member always synchronizes DNS data with the primary server only. It never synchronizes data with the secondary server, even if the primary server fails.

If a Grid member manages several Microsoft servers that are secondary servers of the same zone, it synchronizes DNS data as follows:

- If each Microsoft server is assigned to a different DNS view, the Grid member synchronizes data with each one.
- If the Microsoft servers are synchronized to the same DNS view, the Grid member selects a principal server for synchronization purposes, as follows:
  - The first Microsoft server that is assigned as the DNS secondary server is designated principal server.
  - If the secondary servers are managed in read-only and read/write modes, the Grid member always selects a server that is managed in read/write mode.
  - If a Microsoft server fails three successive synchronizations, it loses its principal server status. The Grid Master checks the date that each server last became a principal server and selects the server that has not been the designated principal server the longest.

Note that a Grid member could fail to synchronize with a Microsoft server due to errors, such as a disabled account or an expired password. In these situations, the failure count is reset and is not increased. This prevents the Microsoft server from losing its master status to another Microsoft server that could experience the same errors.

When a zone is served by multiple Microsoft servers, the **MS Sync Server** column of the **Zones** tab shows which Microsoft server is actually performing the synchronization of that zone with the Grid.

## Enabling and Disabling DNS Zone Synchronization

You can enable and disable synchronization for a DNS zone that is assigned to a Microsoft server. Zone synchronization setting is valid for the selected zone in the DNS view for which it is defined. When you disable the synchronization for a DNS zone, the NIOS appliance stops synchronizing data with the zone but it does not disable the zone. In addition, the zone synchronization setting takes effect regardless of the role of the zone (primary, secondary, or stub), synchronization mode, and permission (read-only or read-write).

Users who have permissions to edit zone properties can enable or disable zone synchronization. Note that you can enable and disable zone synchronization for Microsoft zones only and when you add Microsoft primary or secondary servers.

To configure Microsoft zone synchronization:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *zone* checkbox, and then click the Edit icon.
2. In the *Authoritative Zone* editor, click **Name Servers**.
3. Select **Use this set of name servers**, and enter the details as specified in [Assigning Zone Authority to Name Servers](#).
  - **Disable Zone Synchronization:** Select this to disable zone synchronization for the selected zone. Zone synchronization is enabled when you clear the checkbox.
4. Save the configuration.

## Enabling and Disabling DNS Zone Synchronization for Multiple Zones

You can enable or disable Microsoft zone synchronization for multiple zones:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab, and select the *zone* checkboxes. You can either select a single zone or multiple zones.
2. Expand the Toolbar and click the arrow beside **MS Zone Synchronize** to select an option.
  - **Enable:** Click **Enable** to enable zone synchronization for the selected zones. Click **Yes** in the *Enable MS Zone Synchronization* dialog box to enable synchronization or click **No** to cancel the process.
  - **Disable:** Click **Disable** to disable zone synchronization for the selected zones. Click **Yes** in the *Disable MS Zone Synchronization* dialog box to disable synchronization or click **No** to cancel the process.

## Considerations about Zone Synchronization

Before you enable and disable zone synchronization for a Microsoft DNS zone, consider the following:

- Zone synchronization is enabled by default.
- A zone must be synchronized at least once before you can disable synchronization.
- Zones defined on a Microsoft DNS server are listed in the NIOS appliance, regardless of whether they are synchronized or not.
- When you delete a zone on the Microsoft server, it is automatically deleted from NIOS, irrespective of its synchronization setting.
- When you create a zone on the NIOS appliance and disable zone synchronization, the zone is not created on the Microsoft server until you enable zone synchronization.
- When you disable zone synchronization, you can perform operations for which you have access rights on the NIOS appliance. This includes changes made directly to the zone or as a result of changes made in another zone. Examples: delegation creations and editing name servers.
- When you disable synchronization for a zone and perform certain operations on the respective zone, the outcomes of those operations are not replicated on any Microsoft servers assigned to the zone. For Microsoft primary servers, any resource records or zone properties (including name servers) that you create, modify, or delete on the NIOS appliance are not copied to the Microsoft server. For Microsoft secondary servers and stub servers, any zone properties (including name servers and masters) that you create, modify, or delete on the NIOS appliance are not copied to the Microsoft server.
- When you disable synchronization for a zone, the zone retains the Microsoft server that was last selected as the master before synchronization was disabled. The master retains its role when you enable synchronization again.
- When you disable synchronization, NIOS completes the ongoing process. The synchronization effectively stops at the end of the synchronization. NIOS resumes the synchronization of the zone as soon as the member assigned to the master MS server is notified through Grid replication that the zone is no longer disabled for

synchronization and the zone is overdue for synchronization. This is based on the last time the zone was successfully synchronized and the synchronization interval at the time of re-enabling the synchronization.

- Note that the zones that are disabled for synchronization are not accounted in the overall synchronization status.
- NIOS retains the zone synchronization disable setting when you enable or disable the DNS synchronization setting of any MS server that is assigned to the zone.

## IDN Support for Synchronized DNS Data

Infoblox Grid supports IDNs for synchronized DNS data between the appliance and Microsoft servers. For more information about IDNs, see [Managing Internationalized Domain Names](#).

The appliance stores IDNs in punycode and Microsoft server stores IDNs in \xyz format. Due to this difference at the DNS protocol level, IDNs are not allowed in a zone name when you configure NIOS (primary or secondary) and Microsoft (primary or secondary) servers. For information, see [IDN Support Limitations for Synchronized Data](#) below. If synchronized data between NIOS and Microsoft servers contain IDNs, the IDNs are preserved on the primary server. When a Microsoft server is the secondary server for a zone, MMC (Microsoft Management Console) displays the zone content that contains IDNs in punycode only. Make sure that you use a zone name that complies with the DNS protocol when manually configuring an authoritative zone.

A Microsoft server serves a resource record that contains an IDN in \xyz format when it is configured as the primary server and NIOS as the secondary server. For example, use the \xyz\xyz\xyz.com representation on NIOS for 网络 .com, a zone added on the Microsoft server.

You can add resource records that contain IDNs for the following configurations:

- NIOS is the primary server and Microsoft server is the secondary server: You can add records in IDNs or punycode. The appliance preserves IDNs and does not encode punycode to IDNs. Note that you cannot add a resource record that contains an IDN on the Microsoft server when it is configured as the secondary server.
- Microsoft server is the primary server and NIOS is the secondary server: You can add records on both NIOS and Microsoft servers. You can use IDNs or punycode. You can add IDN records on both servers in this scenario.

The following table summarizes how the servers display resource records that contain IDNs after synchronization:

Primary Server	Input	Secondary Server	Input	NIOS displays records in...	Microsoft server displays records in...
NIOS	Punycode	Microsoft	NA	Punycode	Punycode
NIOS	IDN	Microsoft	NA	IDN	Punycode
Microsoft	Punycode	NIOS	Punycode	Punycode	Punycode
Microsoft	IDN	NIOS	IDN	IDN	IDN

## IDN Support Limitations for Synchronized Data

You cannot configure an authoritative zone and stub zone that contains IDNs for the following configurations:

- When NIOS is the primary server and Microsoft server is the secondary server.
- When Microsoft server is the primary server and NIOS is the secondary server.

## Managing Synchronized DNS Data

When Grid members are configured to manage Microsoft servers in read/write mode, you can use Grid Manager to view, edit and delete the DNS data of those servers. You can add new zones and assign them to a Microsoft server. You can modify the properties of zones synchronized from the Microsoft server and edit their resource records as well. All updates are synchronized to the Microsoft servers at regular intervals.

The following sections provide guidelines for managing the zones and resource records served by Microsoft servers:

- [Adding Zones to Microsoft Servers](#)
- [Setting Zone Properties](#)
- [Deleting and Restoring Synchronized Zones](#)
- [Managing Resource Records in Synchronized Zones](#)

Synchronized zones also support the following features:

- You can import data to zones synchronized with Microsoft servers. Note that the import fails if you try to import unsupported records to a Microsoft zone. For information about importing records, see [Importing Zone Data](#).
- You can copy records to and from zones synchronized with Microsoft servers. When copying records to a Microsoft zone, you can copy only those records that are supported by Microsoft servers. For information about copying records, see as described in [Defining a Match Destinations List](#).

## Adding Zones to Microsoft Servers

From Grid Manager, you can create zones and assign Microsoft servers as their primary or secondary servers. The managing Grid member then synchronizes these zones to the appropriate Microsoft servers.

From Grid Manager, you can add the following types of zones to Microsoft servers:

- Authoritative forward- and reverse-mapping zones — For information, see [Configuring Authoritative Zones](#).
- Forward- and reverse-mapping stub zones — For information, see [Configuring Stub Zones](#).
- Delegations — For information, see [Configuring a Delegation](#).

Note that you cannot add a zone on a Microsoft server and configure it to be served by an Infoblox Grid member. For example, on the Microsoft server, you cannot add a zone and assign a Grid member as its primary server and the Microsoft server as the secondary server. You must add such a zone from Grid Manager.

Following are guidelines for adding zones to a Microsoft server:

- The primary or secondary server of the zone must be a Microsoft server.
- If the primary server is a domain controller, you can enable the option to store the zone in Active Directory, making it an AD-integrated zone. Note that you can enable Active Directory integration only after the Microsoft server has been synchronized at least once because its AD ability is not known before the synchronization.
- You do not have to assign a Grid member as the primary or a secondary server of the zone. For example, a zone can have a Microsoft server as its primary server and an external secondary server.
- The zone must be in the same DNS view to which the DNS data of the Microsoft server was synchronized. You cannot add a zone served by the Microsoft server to a different DNS view.
- The zone does not inherit the properties from the Grid or from the DNS view. It uses the Infoblox-defined defaults. You can change the property values, as described in [Setting Zone Properties](#) below.
- You can set certain zone properties that are not supported and synchronized to the Microsoft server. For example, you can define extensible attributes and administrative permissions. When you set these properties, they apply to the zones only when they are managed from Grid Manager.
- Infoblox does not support all the zone properties of a Microsoft DNS server. When a Grid member synchronizes zones that were created on Grid Manager to the Microsoft server, the zones contain default values for all unsupported properties.
- If you set the Disable option, the zone status is set to "Paused" on the Microsoft server. A zone in a "Paused" status is not served to DNS clients, nor is it available for zone updates.
- Setting the Disable option does not stop synchronization. Grid members synchronize disabled zones.
- The member learns the Windows version of the Microsoft server after its first successful synchronization. Certain zones and resource records are dependent on a specific Windows version. You cannot assign these zones to Microsoft servers whose versions are unknown or insufficient.
- If the member is a secondary server for a zone with a Microsoft primary server, the member obtains the zone data through DNS zone transfers from the Microsoft primary server; not through synchronizations. This ensures that the zone data is always current on the Infoblox secondary server, as it does not have to wait for synchronizations to update its data.

## Setting Zone Properties

When the primary server of a zone is a Microsoft server, it does not inherit its properties from the Grid. Zones that are synchronized from a Microsoft server retain their original properties. Zones that Grid Manager admins create assume the

Infoblox-defined default values.

To modify the properties of a synchronized zone:

1. From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *DNS\_view* -> *zone* checkbox and click the Edit icon.
2. In the *Authoritative Zone* editor, you can do the following in each tab:
  - **General:** You can add or edit comments, and set the Disable and Lock options. Setting the Disable option sets the status of the zone to "Paused" on the Microsoft server. Grid members synchronize disabled zones to Microsoft servers.
  - **Name Servers:** You can modify the name servers assigned to the zone. For information, see [Assigning Zone Authority to Name Servers](#).
  - **Settings:** If the zone was synchronized from a Microsoft server, this tab displays the original settings from the Microsoft server. If the zone was created using Grid Manager, then it inherits the TTL values from the Grid. Note that these values might be different from those on the Microsoft server. To change any of these values, see [Configuring DNS Service Properties](#).
  - **Zone Transfers:** In this tab, you specify the servers to which zone transfers are allowed. For information about zone transfers, see [Enabling Zone Transfers](#). Set the following parameters, depending on whether the primary or secondary servers of the zone are Infoblox or Microsoft DNS servers:
    - If the primary server is an Infoblox, Microsoft or external primary and the secondary servers are both Infoblox and Microsoft DNS servers, this tab displays two separate tables where you can specify zone transfer settings for the Infoblox DNS servers and the Microsoft DNS servers.  
**Zone Transfer Settings for Infoblox Members:** Specify the settings as described in [Configuring Zone Transfers](#).  
**Zone Transfer Settings for Microsoft Servers:** Note that you cannot use a named ACL for access control though you can use individual ACEs. For information about named ACLs and access control, see [Configuring Access Control](#). You can set access control for zone transfers for Microsoft servers to one of the following:
      - **None:** Does not allow zone transfers to any name server.
      - **Any:** Allows zone transfers to any IP address.
      - **Any Name Server:** Allows zone transfers to any name server in the Name Servers table.
      - **Address:** Allows zone transfers to the IP address that you specify.
    - If both the primary and secondary servers are Microsoft servers, the dialog box displays the **Zone Transfer Settings for Microsoft Servers** table only.
    - If no Microsoft servers are primary or secondary servers, then the dialog box displays the **Zone Transfer Settings for Infoblox Members** table only.
  - **Updates:** In this tab, you specify whether the zone can accept dynamic DNS updates. For information about dynamic DNS updates, see [Configuring DDNS Updates](#). If the primary server is a Microsoft server, regardless of the secondary servers, the **Updates** tab displays the following:
    - **Dynamic Updates:** Select one of the following:
      - **None:** The zone does not accept dynamic updates.
      - **Secure Only:** This appears only if the zone is AD-integrated. The zone accepts GSS-TSIG-signed updates only.
      - **Nonsecure and Secure:** The zone accepts both nonsecure and GSS-TSIG-signed updates.
    - **Active Directory:**
      - **Automatically create underscore zones:** This option allows the appliance to create the following subzones that the DNS server must have to answer AD-related DNS queries:
        - `_msdcs.zone`
        - `_sites.zone`
        - `_tcp.zone`
        - `_udp.zone`
        - `domaindnszones.zone`
        - `forestdnszones.zone`

Note that these zones are automatically generated. You cannot edit these zones or import data into them. They cannot be modified, thus providing protection against forged updates.

- **Extensible Attributes:** Extensible attributes apply to the zones only when they are managed from Grid Manager. For information, see [Using Extensible Attributes](#).



- **Permissions:** These permissions apply to Infoblox Grid Manager administrators only. For information, see [About Administrative Permissions](#).

## Deleting and Restoring Synchronized Zones

When you delete a synchronized zone from the Grid, Grid Manager moves the zone and its resource records to the Recycle Bin. It deletes the zone and its resource records from the Microsoft server at the next synchronization. Note that if you delete a zone on Grid Manager and plan to add it back to the database with different properties or resource records, ensure that you wait until after the deletion is synchronized to the Microsoft server to add the new zone. Otherwise, if you delete a zone and add a new zone with the same name within a synchronization interval, Grid Manager will synchronize the zone properties and resource records from the Microsoft server to the newly added zone on Grid Manager.

If a zone has subzones, you can choose to remove them and their resource records or "reparent" them to the parent zone of the one you are removing. For information, see [Removing Zones](#).

If you restore deleted zones from the Recycle Bin, the Grid member restores it on the Microsoft server as well. For information, see [Restoring Zone Data](#).

## Managing Resource Records in Synchronized Zones

From Grid Manager, you can add and edit resource records in zones served by Microsoft servers. For information about adding and managing resource records, see [Managing Resource Records](#). You can also use IP Map and the IP List to track A, AAAA and PTR records that are synchronized from Microsoft servers. For information, see [About IP Address Management](#).

Microsoft DNS servers support all the resource records supported by Infoblox DNS servers, except for hosts, bulk hosts and shared record groups. You cannot add these records to zones served by Microsoft servers or assign zones with these records to Microsoft servers.

Following are guidelines for adding and managing resource records in synchronized zones:

- Infoblox DNS servers support defining a naming policy for the hostnames of A, AAAA, MX, and NS records based on user-defined or default patterns. For information, see [Specifying Hostname Policies](#). The hostname policy applies only when records are created from Grid Manager. Resource records that originate from the Microsoft server are synchronized to the Grid member even if they do not comply with the hostname policy of the Grid member. The policy is enforced only if you edit the resource record from Grid manager.
- When you create an A or AAAA resource record on the NIOS appliance with the option to automatically create the corresponding PTR record, Grid Manager uses the deepest reverse zone that can hold the record. For example, a Grid has the following reverse zones: 10.in-addr.arpa, 0.10.in-addr.arpa, and 0.0.10.in-addr.arpa. When you create the A record www A with the IP address 10.0.0.1, Grid Manager creates a PTR record in the zone 0.0.10.in-addr.arpa. If the deepest zone does not allow the creation of the PTR record, Grid Manager creates the A record, but not the PTR record, and displays a warning.

Note that if the **Enable PTR record removal for A/AAAA records** option is selected and if you try to delete the A or AAAA records in zones served by Microsoft servers, the appliance displays the *Delete Confirmation (A or AAAA Record)* dialog box to confirm whether you want to remove the corresponding PTR record that was automatically generated while creating the A or AAAA record. In the *Delete Confirmation* dialog box, select the **Remove associated PTR resource record(s)** checkbox and click **Yes** to delete the associated PTR record or click **No** to cancel. For information about enabling this option, see [Deleting PTR Records associated with A or AAAA Records](#).

- You can add and edit DNAME records in a DNS zone assigned to a Microsoft server running Windows 2008 or Windows Server 2012. You cannot add or edit DNAME records in zones assigned to Microsoft servers running earlier Windows versions.
- You can disable synchronized resource records from Grid Manager. When you disable a resource record, it is removed from the Microsoft server at the next synchronization.
- If you add a resource record with invalid data from Grid Manager, such as a DNAME record with an alias name that has special characters, the invalid resource record is not synchronized to the Microsoft server and is eventually deleted from the Grid. The error is logged in the Microsoft log.
- If the zone of the resource record was created using Grid Manager, then it and all its resource records inherit their TTL values from the Grid. Note that these values might be different from those on the Microsoft server. You can

change these values to match those on the Microsoft server. For information on changing these values, see [Configuring DNS Service Properties](#).

- Grid Manager and Microsoft DNS servers display TXT records differently.

On Grid Manager, you enter the text string of TXT records as defined in RFC 1035. You can enter the following:

- A contiguous set of characters without spaces. If you enclose the characters in double quotes, Grid Manager displays the character string without the double quotes. For example, if you enter **"abcdef"**, Grid Manager displays **abcdef**.
- A string that contains any character, including spaces, enclosed in quotes.
  - If the string contains a quote ("), you must precede it with a backslash.
  - If you enter a text string with multiple spaces between each word and the string is not enclosed in double quotes, Grid Manager displays the text string with a single space between each word. For example, if you enter **text string**, the GUI displays **text string**. To preserve multiple spaces, enclose the string in double quotes.

Unlike on Microsoft DNS servers, you cannot enter a text string on multiple lines in Grid Manager. However, each contiguous set of characters or quoted string entered on Grid Manager is equivalent to a separate line entered on a Microsoft DNS server.

On a Microsoft DNS server, you can enter text without quotes and with each line on a separate line. Microsoft DNS servers then display the text in a brief format where the lines are separated by a comma and a space.

For example, if you enter the following in the **Text** field of the TXT Record wizard or editor on Grid Manager:

**"this is a line""with another line""and a third one"**

It is served by the Microsoft and Infoblox DNS servers as:

**"this is a line""with another line""and a third one"**

But it is displayed in the Microsoft DNS server as:

**this is a line, with another line, and a third one**

## Synchronizing Updates

A Grid member synchronizes DNS data with each managed Microsoft server at regular intervals. Grid Manager admins with the applicable permissions can then update the synchronized DNS zones and resource records. During each synchronization, updates from Grid Manager are applied to the Microsoft server and updates from the Microsoft server are applied to the Grid as well. Note that the resource records are synchronized only if there is a change to the SOA record on either the Microsoft server or the Grid.

The following examples illustrate how Grid members synchronize DNS data:

- If a Microsoft server admin adds the `finance.corpxyz.com` zone, it is also added to the Grid after a synchronization.
- If a Grid Manager admin changes the A record of `admin.corpxyz.com` from `10.2.1.5` to `10.2.1.6`, the IP address of its corresponding A record on the Microsoft server is updated to `10.2.1.6`.
- If a Grid Manager admin deletes a DNS zone that is assigned to a Microsoft server, the corresponding zone on the Microsoft server is deleted as well in the next synchronization.

Because admins can update DNS data from the Microsoft server and from Grid Manager, conflicts can occur during synchronization. In addition, Microsoft servers and Infoblox DNS servers have some differences in the features they support and the way they handle certain zones and resource records.

The following guidelines describe how the Grid member resolves conflicts and handles any differences when DNS data is synchronized between a Microsoft server and the Grid.

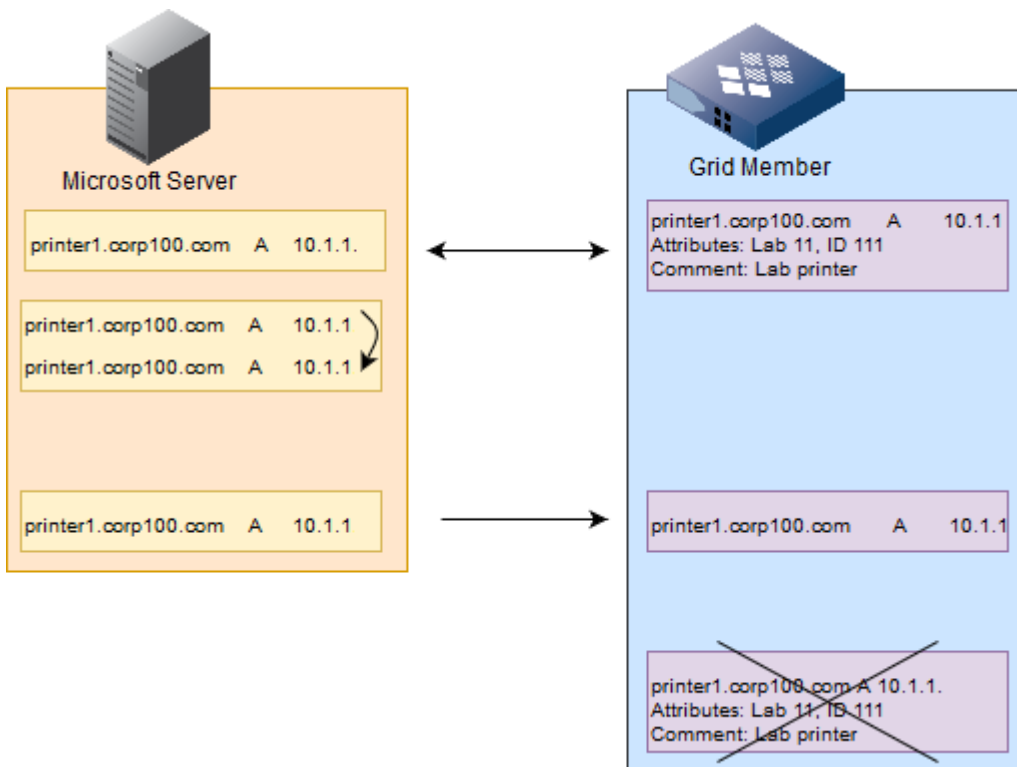
- On Microsoft servers, users can enter FQDNs and labels using a mix of upper and lower case characters. The servers preserve the original case when they store the data. When the Grid member synchronizes data with the Microsoft server, it displays the data in lower case in Grid Manager and the Infoblox API. The case of the data is preserved as long as no change is made to the DNS zone or resource record. If a Grid Manager admin modifies a DNS zone or resource record, the next synchronization converts the object name to lower case on the Microsoft server.
- If a Microsoft server admin modifies an object that has a pending scheduled task and synchronization occurs before the scheduled task, the object is modified in both the Microsoft server and the Grid member. When the scheduled task executes at its scheduled time, it fails and an error message is logged in the audit log.

- A situation could arise where two Microsoft servers in different domains are primary servers for zones with the same name. For example, two reverse-mapping zones could be named 1.1.10.in-addr-arpa in two Microsoft servers managed by the same member. If the two Microsoft servers are synchronized to different DNS views, the Grid member synchronizes each one separately. If the Microsoft servers are synchronized to the same DNS view, then the Grid member synchronizes the zone with the first Microsoft server. During the synchronization with the second Microsoft server, the Grid member logs an error and does not synchronize the zone.
- The Grid member does not synchronize the naming policy configured on Microsoft servers. Zones and resource records that fail the policy check on Microsoft servers are reported in the synchronization log file.
- When you remove a Microsoft server that is assigned to a zone, the succeeding synchronization removes the zone from the Microsoft server.
- When a Microsoft server admin and a Grid Manager admin change the same object, the Grid member retains the version that exists on the Microsoft server. Following are some examples:

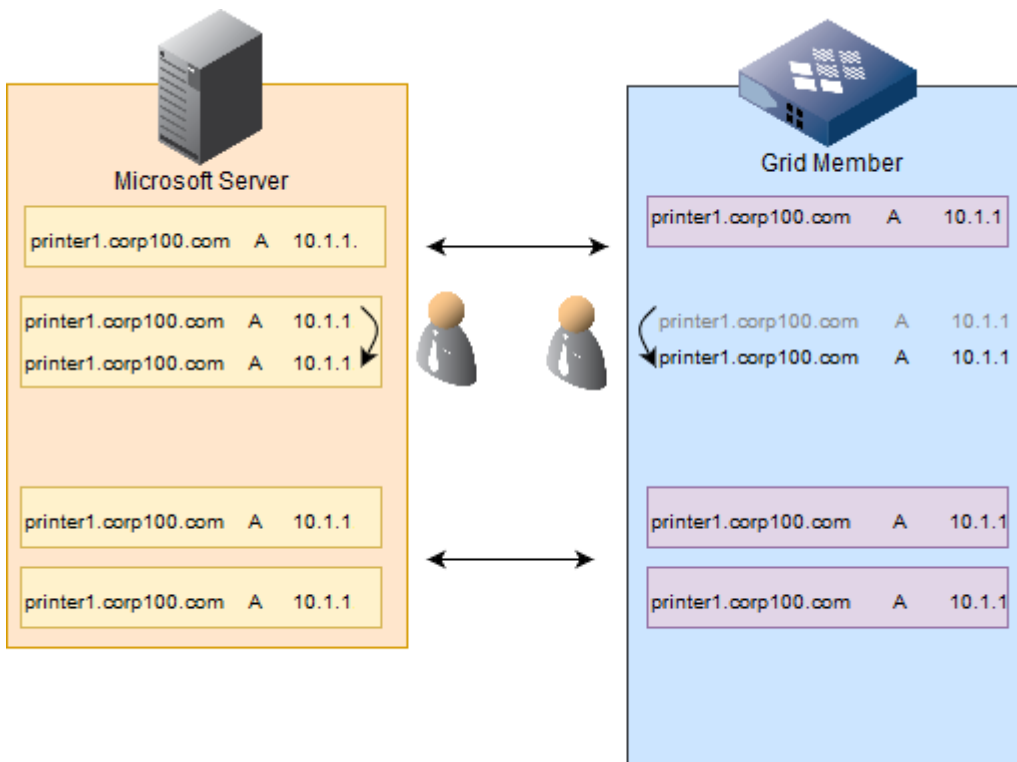
Grid Manager Admin...	Microsoft Server Admin...	After Synchronization
Deletes the corpxyz.com zone	Updates the corpxyz.com zone	The corpxyz.com zone is created on the Grid with the updates and is assigned to the Microsoft server .
Changes the zone transfer settings of the sales.corpxyz.com zone.	Deletes the sales.corpxyz.com zone.	The sales.corpxyz.com is deleted from the Grid as well.

- Changing the name or IP address of a resource record on the Microsoft server effectively deletes the original resource record and creates a new record with the current information. During the synchronization, the Grid member also deletes the original record, including its associated properties, such as its extensible attributes and administrative permissions, and creates a new record.

For example, as shown in the below figure, the A record for printer1.corpxyz.com is on both the Microsoft and Infoblox Grid member. On the Grid, the A record has extensible attributes and a comment. A Microsoft server admin changes the IP address of the A1 resource record from 10.1.1.2 to 10.1.1.3. On the Microsoft server, this is equivalent to deleting the A1 resource record with the IP address 10.1.1.2 and then adding a new A1 resource record with the IP address 10.1.1.3. When the data is synchronized, the Grid member deletes the original record with its extensible attributes and comments and creates a new A record with IP address 10.1.1.3.



- If a Microsoft server admin changes the IP address of a resource record and a Grid Manager admin changes the IP address of the same resource record, they are effectively deleting the record and each creating a new one. For example, as shown in the below figure, a Microsoft server admin changes the IP address of the A resource record for printer1.corpxyz.com from 10.1.1.2 to 10.1.1.3, and a Grid Manager admin changes the IP address of the same resource record to 10.1.1.4. When the data is synchronized, the Grid member deletes the A1 resource record with IP address 10.1.1.2 and creates an A resource record with IP address 10.1.1.3 and another A1 resource record with IP address 10.1.1.4.



- The Microsoft server does not allow the creation of arpa subzones as forward-mapping zones, similarly, the appliance restricts assigning arpa subzones (zone names ending with .arpa) to the Microsoft server.
- NIOS does not synchronize the top-level reverse-mapping zones (in-addr.arpa and ip6.arpa) created on the Microsoft server and the top-level reverse-mapping zones (in-addr.arpa and ip6.arpa) created on the NIOS appliance cannot be assigned to the Microsoft server.
- Grid members can synchronize classless IPv4 reverse-mapping zones from the Microsoft server to the Grid only if the zone prefix is in one of the following formats: `<subnet>/<subnet mask bit count>` or `<subnet>-<subnet mask bit count>`. For example, `128/26.2.0.192.in-addr.arpa`. If the zone prefix is not in the specified format, the Grid member skips the zone and logs an error message. For information, see <http://technet.microsoft.com/en-us/library/cc961414.aspx>.  
Likewise, Grid Manager admins can add a classless IPv4 reverse-mapping authoritative or stub zone to a Microsoft server only if its prefix is in the specified format. For information about configuring classless IPv4 reverse-mapping zones in Grid Manager, see [Specifying an RFC 2317 Prefix](#).
- Grid members synchronize DNS records that contain values that Infoblox does not support. Grid Manager admins can view these records, but they cannot edit or restore such records. For example, if a member synchronizes a NAPTR record that contains an unsupported value in the Service field, admins can view this record but they cannot edit or restore it, as long as it contains an unsupported value.
- When a Grid member synchronizes a zone from a Microsoft server to the appliance and that zone contains UTF-8 characters in the "Responsible Person" field, Grid Manager displays the "Responsible Person" value in the RNAME field of the SOA record of the zone. Note though that you cannot edit the SOA record if the RNAME field contains unsupported UTF-8 characters.
- Synchronizing a new zone from the Microsoft server to the Grid is a two-step process. First the zone name is synchronized, and then its properties and records are synchronized. The zone synchronization from the Microsoft server is not considered complete until both steps are done. NIOS drops any records that are created on the appliance for the synchronized zone before it is completely synchronized.  
For example, the `corpxyz.com` zone is created on a Microsoft server and then synchronized to the NIOS appliance. If a NIOS admin creates a record, such as an A record, in the `corpxyz.com` zone before it is fully synchronized, the record is removed from the `corpxyz.com` zone. Ensure that both the zone and its contents are completely synchronized before you add a record to a zone on the NIOS appliance.

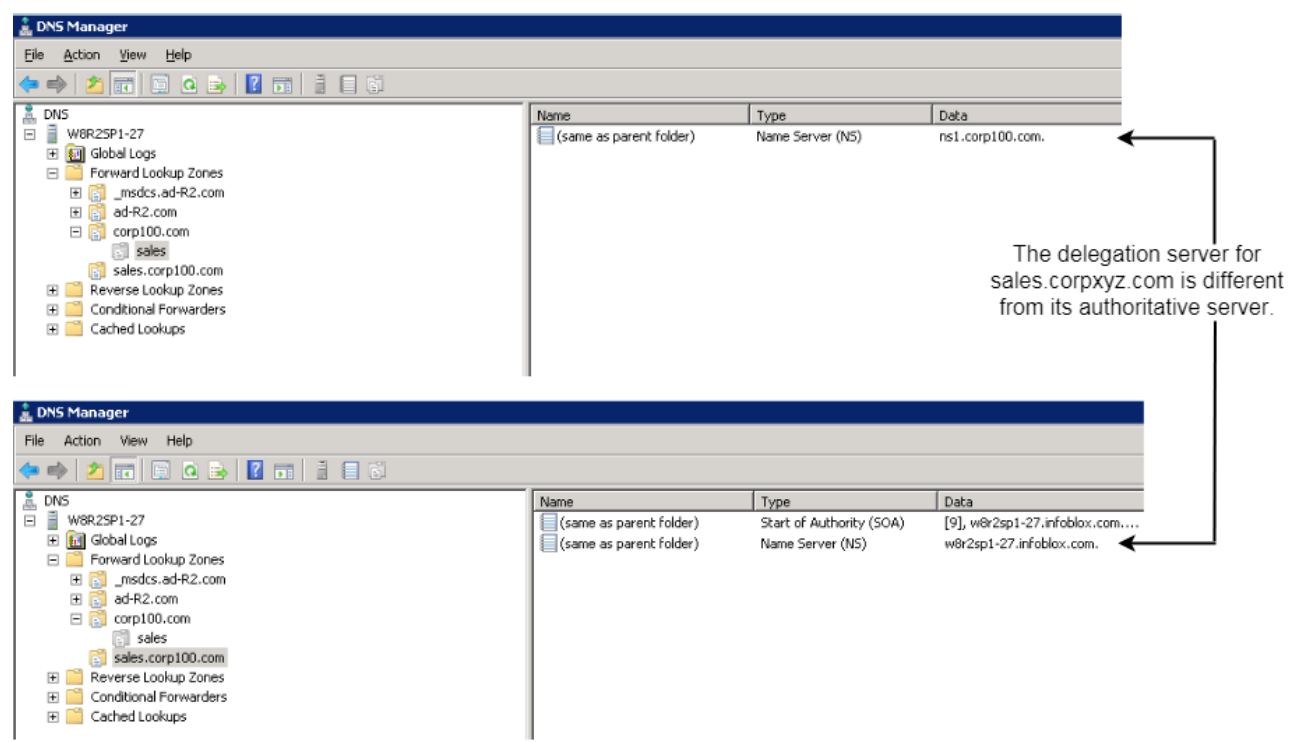
- When a Microsoft admin deletes a zone whose primary and secondary servers are Microsoft servers, the zone is deleted from the Grid after a synchronization. If the secondary Microsoft servers are managed by the member in Read-Write mode, the zone is removed from the secondary servers as well. But if some of the secondary Microsoft servers are managed by the member in Read-Only mode, then at the next synchronization the zone is recreated on the Grid with the Microsoft servers as the secondary servers and the masters defined for the zone as external primary servers.
- If you add Grid members or other Microsoft servers as secondary servers to a zone on the Microsoft server, NIOS does not automatically add them as Grid Secondary or Microsoft Secondary servers in the Name Servers tab of the zone after the synchronization. Instead, NIOS creates NS records for them in the zone.

### Synchronizing Delegations

When a parent zone delegates a subdomain to one or more name servers, Infoblox DNS servers require the delegation name servers to also be authoritative for the subzone. Microsoft servers do not; they allow the delegation servers of a subzone to be different from its authoritative servers. Infoblox DNS servers support this configuration only if the primary server of the parent zone is a Microsoft server. This configuration is retained when delegations are synchronized from Microsoft servers to the Grid.

For example, as shown in the below figure, on a Microsoft server, corpxyz.com delegates sales.corpxyz.com to the name server ns1.corpxyz.com; but the authoritative server of sales.corpxyz.com is 2k3r264-2.infoblox.com.

*Delegation Server and Authoritative Server for corpxyz.com*



The following figure shows that after corpxyz.com and its subzone are synchronized to the Grid, corpxyz.com contains an NS record for sales.corpxyz.com and an A record for the delegation name server ns1.corpxyz.com. The *MS Delegation Addresses* column displays the IP address of the delegation server of the subzone sales.corpxyz.com.

*corpxyz.com Synchronized to the Grid*

corp100.com Authoritative Zone

Records Subzones

Quick Filter: None Filter On Show Filter Toggle flat view

NAME	TYPE	DATA	MS DELEGATION ADDRESSES	COMMENT
	NS Record	w8r2sp1-27.ad-r2.com		Auto-created by Add Zone
	SOA Record	Serial: 83 MNAME: w8r2sp1-27.ad- RNAME: hostmaster@ad Refresh: 900 Retry: 600 Expire: 86400 Negative Caching TTL: 3600		Auto-created by Add Zone
ns1	A Record	10.2.3.4		
sales	NS Record	ns1.corp100.com	10.2.3.4	Auto-created by Add Zone

After the synchronization, you can add name servers for the delegation as follows:

1. Select the zone by navigating to the **Data Management** tab -> **DNS** -> **Zones** -> *parent\_zone*.
2. Click the Add icon and select **Record** -> **NS Record**.
3. Complete the following and click **Next**:
  - **Name Server**: Enter the hostname you want to configure as the name server for the zone.
  - **Name**: Specify the name of the subzone. Note that you cannot change this name when you edit the record.
4. Enter the IP address of the name server.
5. Save the configuration.

NIOS adds an NS record for the new delegation server and synchronizes this update to the Microsoft server. In the following figure, a new delegation server, ns2.corpxyz.com, was added.

*NS Record for ns2.corpxyz.com*

corp100.com Authoritative Zone

Records Subzones

Quick Filter: None Filter On Show Filter Toggle flat view

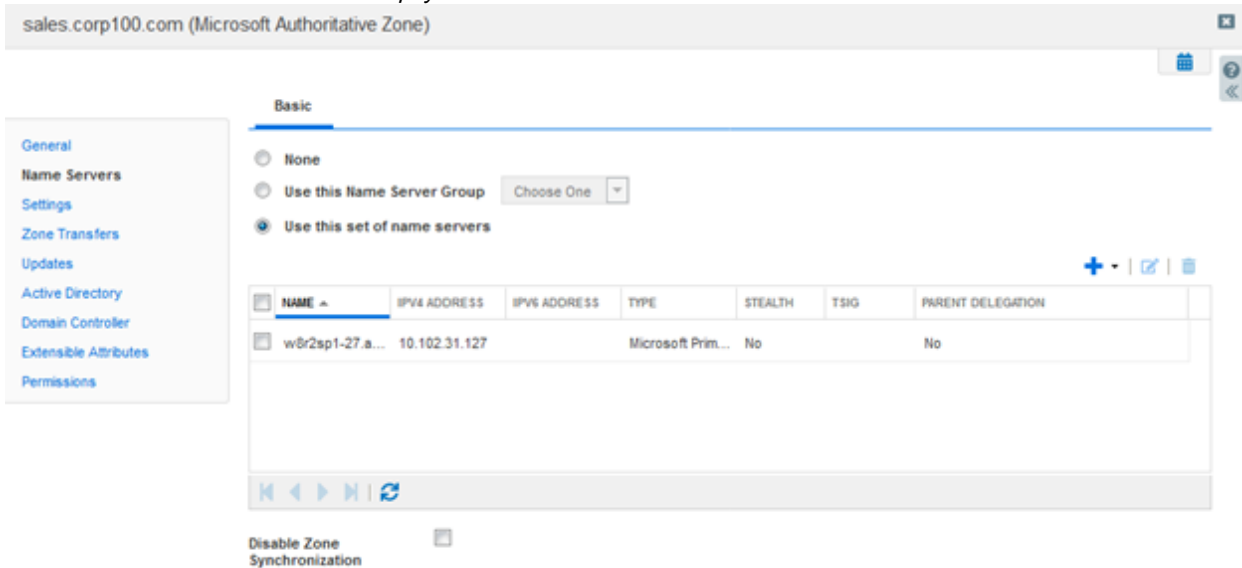
NAME	TYPE	DATA	MS DELEGATION ADDRESSES	COMMENT
	NS Record	w8r2sp1-27.ad-r2.com		Auto-created by Add Zone
	SOA Record	Serial: 67 MNAME: w8r2sp1-27.ad- RNAME: hostmaster@ad Refresh: 900 Retry: 600 Expire: 86400 Negative Caching TTL: 3600		Auto-created by Add Zone
ns1	A Record	10.2.3.4		Auto-created by Add Zone
ns2	A Record	10.1.2.3		Auto-created by Add NS
sales	NS Record	ns2.corp100.com	10.1.2.3	Auto-created by Add Zone
sales	NS Record	ns1.corp100.com	10.2.3.4	Auto-created by Add Zone

When you navigate to the **Name Servers** tab of sales.corpxyz.com to view the authoritative name servers for the



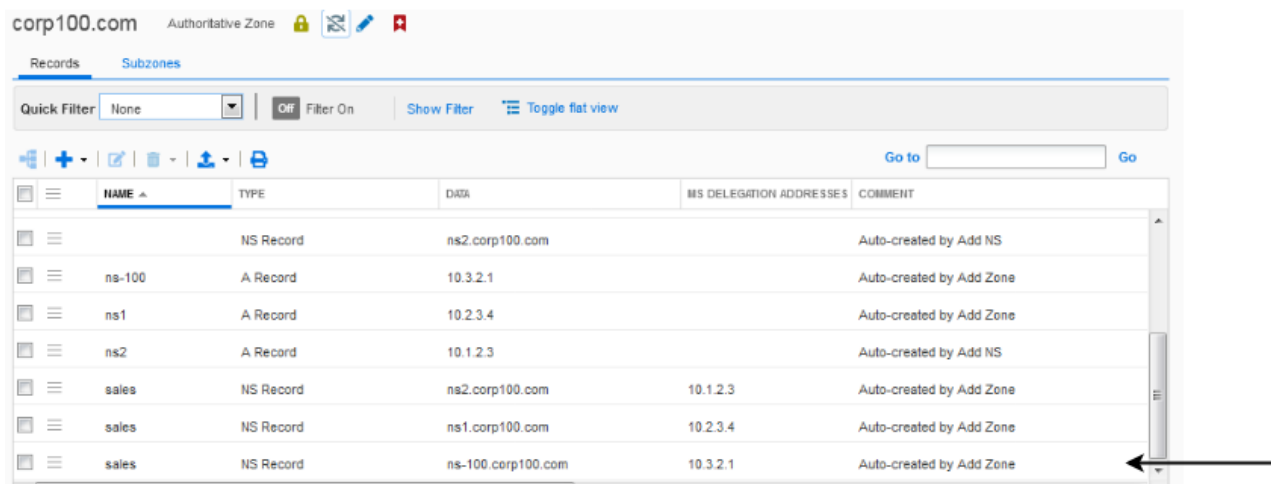
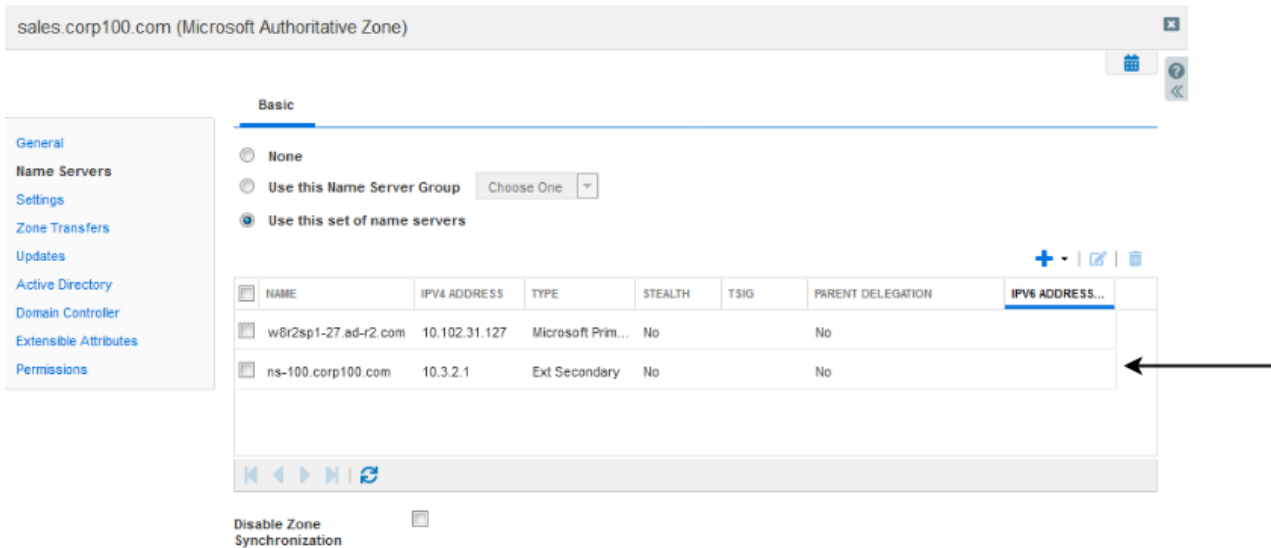
subzone, note that as shown in the following figure, the table displays 2k8r264-2.infoblox.com as the authoritative server for the subzone. The **Parent Delegation** column indicates if the FQDN and IP address of the authoritative name server for the zone matches the FQDN and IP address in the delegation zone's NS record. In the example, the authoritative name server 2k8r264-2.infoblox.com is different from the delegation name servers (ns1.corpxyz.com and ns2.corpxyz.com), so the column displays **No**.

*Authoritative Name Server of sales.corpxyz.com*



Note though that because Infoblox DNS servers require the delegation servers to also be authoritative for the subzone, if you add another authoritative name server to the subzone from Grid Manager, NIOS also adds it as a delegation server in the parent zone. For example, as shown in the below figure, when an admin adds the name server ns-100.corpxyz.com as an external secondary server for sales.corpxyz.com, NIOS automatically adds it as a delegation server by adding an NS record for it in the parent zone.

*Adding Another Authoritative Server from Grid Manager*



## Synchronizing AD-Integrated Zones

An AD-integrated zone can be served by multiple domain controllers, and a Grid member can manage more than one of the domain controllers. If the domain controllers are configured to synchronize their DNS data to different DNS views, the Grid member synchronizes DNS data with each domain controller. If the domain controllers are configured to synchronize their DNS data to the same DNS view, the member selects a principal server for synchronization purposes and synchronizes data with that principal server only. The selection of the principal server is logged, as well as when it changes. The Grid member selects a principal server as follows:

- The first domain controller that is assigned as the primary server is designated principal server.
- If a domain controller fails three successive synchronizations, it loses its principal status. The Grid Master then checks the date that each domain controller last became a principal server and selects the one that has not been the designated principal the longest.
- If the domain controllers are managed in read-only and read/write modes, the Grid member always selects the domain controller that is managed in read/write mode.

When a zone is served by multiple Microsoft servers, the **MS Sync Server** column of the **Zones** tab shows which Microsoft server is actually performing the synchronization of that zone with the Grid.

The Grid member periodically checks if each zone has one principal server. If it does not find a principal server for a zone, the Grid member selects one among the name servers assigned to the zone. It gives priority to the server that was not the designated principal server the longest.

Following are additional guidelines when synchronizing AD-integrated zones:

- You can create an AD-integrated zone on Grid Manager and assign one domain controller as its primary server. If a domain controller admin adds more primary servers to the zone, they are added to the zone on Grid Manager when the zone is synchronized. If you want to delete the primary servers, you must delete all the primary servers at once. You cannot delete only a subset of the servers.
- A situation could arise where two domain controllers in different AD domains are primary servers for zones with the same name. For example, two reverse-mapping zones could be named 1.1.10.in-addr-arpa in two domain controllers managed by the same member. If the two domain controllers are synchronized to different DNS views, the Grid member synchronizes each one separately. If the domain controllers are synchronized to the same DNS view, then the Grid member synchronizes the zone with the first domain controller. During the synchronization with the second domain controller, the Grid member logs an error and does not synchronize the zone.
- If a Grid Manager admin deletes a CNAME record that has a blank canonical name from an AD-integrated zone, this CNAME record is not deleted from the Microsoft server after the synchronization if the AD-integrated zone is hosted on a Microsoft server running Windows 2008 R2 or Windows Server 2012.
- When a Microsoft server is the primary server of a zone that contains an `_msdcs` zone, it appends the parent zone name to the server name in the NS record of the `_msdcs` zone. But when an Infoblox Grid member is the primary server of a zone that contains an `_msdcs` zone, it specifies the server name only in the NS record. For example, the `_msdcs` zone is in the `corpxyz.com` zone and the name server is `nameserver100.com`. When a Microsoft server is the primary server of `corpxyz.com`, the server name on the NS record of the `_msdcs` zone is `nameserver100.com.corpxyz.com`. When a Grid member is the primary server, the server name on the NS record of the `_msdcs` zone is `nameserver100.com`.

## Resolving Conflicts

Some conflicts require intervention from an admin. For example, a Grid member cannot synchronize a zone when its primary server on the Microsoft server is different from its primary server on the Grid. When a Grid member is unable to synchronize data due to such conflicts, it logs an error, skips the object with the error and continues synchronizing the rest of the data. You can then view the Microsoft logs to check which objects were not synchronized. If you resolve the problem, the Grid member synchronizes the object on its next attempt. For information about the logs, see [Viewing Synchronization Logs](#).

## Viewing Members and Managed Servers

You can view Infoblox and Microsoft DNS servers by navigating to the **Data Management** tab -> **DNS** tab, and then selecting the **Members/Servers** tab. The panel displays the following information about each DNS server:

- **Name:** The hostname of the Grid member or Microsoft server.
- **Status:** The status of the DNS service on the Grid member or Microsoft server.
- **Comment:** Comments that were entered for the Grid member or Microsoft server.
- **Site:** Values that were entered for this pre-defined attribute.
- **Address:** The IP address of the Grid member or Microsoft server. You can do the following:
  - List the DNS views or zones served by the member or Microsoft server.
    - Click a Grid member or Microsoft server name.
  - Edit the properties of a Grid member or Microsoft server.
    - Click the checkbox beside a Grid member or Microsoft server, and then click the Edit icon. To edit the DNS properties of a Grid member, see [Configuring DNS Service Properties](#).  
To edit the DNS properties of a Microsoft server, see [Specifying Forwarders for Microsoft Servers](#).
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Export the list of Grid members and Microsoft servers to a .csv file.
  - Click the Export icon.

- Print the list of Grid members and Microsoft servers.
  - Click the Print icon.

## Specifying Forwarders for Microsoft Servers

A forwarder is a name server to which all other name servers first send queries that they cannot resolve locally. The forwarder then sends these queries to DNS servers that are external to the network, avoiding the need for the other name servers in your network to send queries off-site. You can define a list of forwarders for each managed Microsoft server as follows:

1. From the **Data Management** tab, select the **DNS** tab -> **Members/Servers** tab -> *ms\_server* checkbox -> Edit icon.
2. Click the Add icon and enter the IP address of the forwarder in the text field.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Disabling and Removing Microsoft DNS Servers

When you disable synchronization with a Microsoft server, the managing Grid member completes any ongoing synchronization and does not start a new one. Synchronization resumes when the Microsoft server is re-enabled. The synchronized DNS data stays in the same state until synchronization resumes. For information, see [Disabling Synchronization](#).

When you remove a managed Microsoft server from the Grid, the managing Grid member terminates any ongoing synchronization and does not start a new one. Zones and their content on the Microsoft server remain in the state that existed the moment the Microsoft server was removed. The Grid retains the zones that were assigned to the Microsoft server that was removed, but deletes the Microsoft server from its assigned zones as follows:

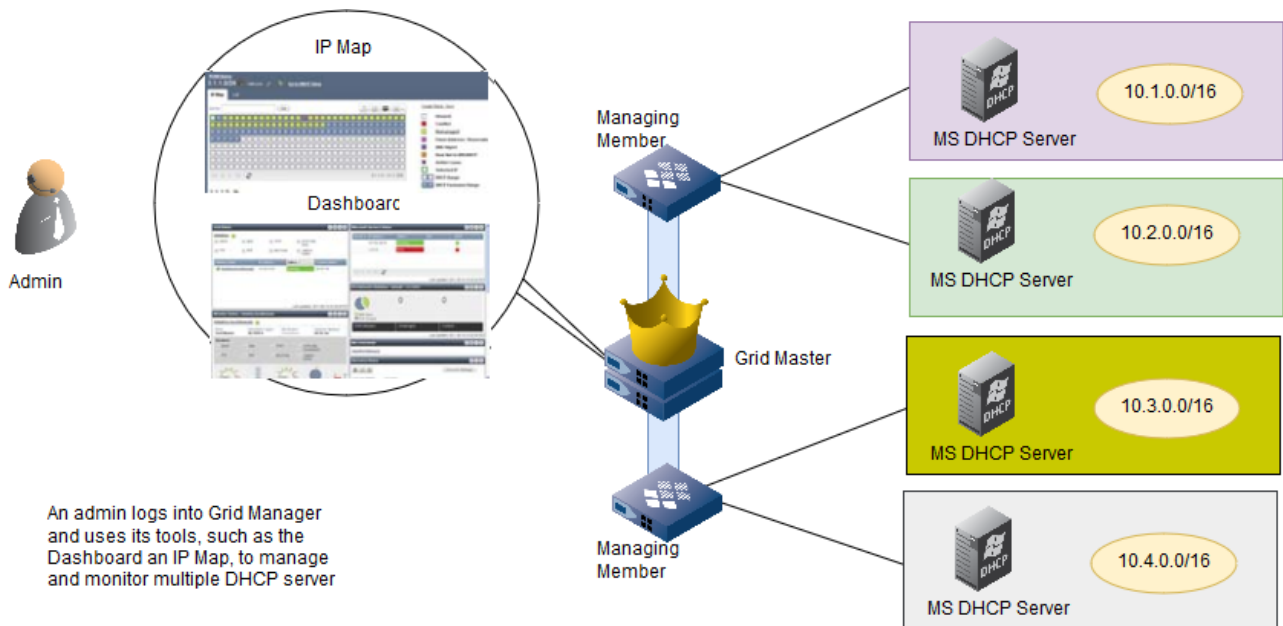
- If the Microsoft server is the only primary server and there are no other assigned servers or if the secondary server is an external secondary server, Grid Manager deletes all the server assignments.
- If the Microsoft server is the only primary server and there are Grid secondary servers, an external primary is created with the FQDN and IP address of the removed Microsoft server.
- If the Microsoft server is a secondary server and there is a Grid primary, an external secondary is created with the FQDN and IP address of the removed Microsoft server.
- If the Microsoft server is a server for a stub zone, the server is removed. To remove a Microsoft DNS server:

1. From the **Data Management** tab, select the **DNS** tab -> **Members/Servers** tab -> *ms\_server* checkbox.
2. Expand the Toolbar and click **Delete**.
3. Click **Yes** when the confirmation dialog box appears.

## Managing Microsoft DHCP Services

Grid Manager enables you to centrally manage the DHCP data of multiple Microsoft DHCP servers from a single interface. Once the DHCP data is synchronized, you can use the Dashboard on Grid Manager to monitor DHCP and server operations, or organize DHCP data into Smart Folders. Through IPAM tools, such as network maps and IP maps, you can track and manage IP address usage in your networks and monitor DHCP range utilization. You can also run a network discovery to retrieve IP allocation for both managed and unmanaged devices— including virtualized resources. For information about the IPAM features, see [About IP Address Management](#).

*Managing Microsoft DHCP Servers from Grid Manager*



This section provides guidelines for using Grid Manager to manage Microsoft DHCP servers and for synchronizing DHCP data between Microsoft servers and the Grid. It discusses features of only Microsoft DHCP servers as they relate to the synchronization of data. Please review the Microsoft documentation for complete information about Microsoft DHCP servers and their features.

In addition, if you encounter technical issues with your Microsoft DHCP servers, contact Microsoft Technical Support or consult the Microsoft Support site at <http://support.microsoft.com/>. Some Windows versions require certain updates and hotfixes installed, so the Microsoft server can synchronize with the Grid member. For information about these requirements, see [Requirements](#).

The topics in this section include:

- [Viewing Synchronized Leases](#)
- [Synchronizing DHCP Failover Relationships](#)
- [Synchronizing IP Addresses with Invalid MAC Addresses](#)
- [Managing Synchronized DHCP Data](#)
- [Synchronizing DHCP Data from Microsoft Servers](#)
- [Managing Microsoft DHCP Servers](#)

## Viewing Synchronized Leases

A Grid member synchronizes all leases from its managed Microsoft server to the Grid. Microsoft servers automatically generate a static lease for each reservation. These static leases are synchronized to the Grid as well. You can view the synchronized leases by navigating to the **Data Management** -> **DHCP** -> **Leases** tab. For information about viewing current leases, see [Viewing Current Leases](#). You can do the following:

- View lease details, by selecting a lease and clicking the Lease Details icon. For additional information, see [Viewing Detailed Lease Information](#).
- Clear a lease, by selecting it and clicking the Clear Lease icon. Note that Grid Manager clears the lease immediately. It does not wait for the next synchronization.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information about using quick filters, see [Finding and Restoring Data](#).

Additionally, you can enable a Grid member to log lease related operations, so you can view these events in the Lease History panel. For information about configuring the lease logging member, see [Configuring DHCP Logging](#) and see [Viewing Lease History](#).

## Synchronizing DHCP Failover Relationships

The synchronization of DHCP data with Microsoft DHCP servers running Microsoft Windows version 2012 or later includes the synchronization of DHCP failover relationships and the associated scopes. Note that this feature does not have any impact on the synchronized data with Microsoft servers running a Windows version earlier than 2012. When you change the synchronization mode of a Microsoft server, it affects the way a failover relationship is synchronized and replicated, but you cannot change these settings directly. When a failover relationship synchronization mode changes from read-only to read/write, NIOS copies the DHCP configuration within the relationship between both partners, using the primary Microsoft server as the reference. When a failover relationship synchronization mode changes from read/write to read-only, NIOS does not change the DHCP configuration within the relationship on each Microsoft server.

### About Microsoft DHCP Failover Relationships

A failover relationship represents the relationship between two Microsoft DHCP servers, where each of them is called a partner. A failover relationship that has at least one Grid member as a partner is called a NIOS failover relationship and if it has at least one Microsoft server as a partner, it is defined as a Microsoft failover relationship. A failover relationship can be configured with two Microsoft servers, a Microsoft server and an external server, two Grid members, a Grid member and an external server.

When you synchronize a failover relationship from a Microsoft server, the NIOS appliance copies the changes originating from the Microsoft server to the failover relationship of the partner, only if the two Microsoft servers match. Two failover relationships match if they have the same name, and if the address of the partner defined in one Microsoft server matches with the address of the other, and vice versa. Also, you must configure Microsoft servers to manage DHCP and enable synchronization on both the Microsoft servers.

The changes that originate from the Microsoft server are applied to the failover relationship on NIOS, and vice versa. The changes are effective on the partner server from its next scheduled synchronization.

The failover relationship synchronization mode for Microsoft is read/write if both partners are each Microsoft servers that are enabled for DHCP synchronization. For all other cases, the failover relationship synchronization is read only.

When NIOS fails to match a failover relationship, the partner of that relationship is considered to be unknown or external. In the case of multiple servers, you must define the Microsoft server on the appliance using the same IP and/or FQDN. Infoblox recommends that you do the following for Microsoft servers in a failover relationship:

- The synchronization interval of two Microsoft servers in a failover relationship must be identical and small.
- Use the same managing member for Microsoft servers in a failover relationship.
- After a Microsoft server fails, you must turn off NIOS synchronization during restore to avoid transferring the old configuration to its partners.
- The primary and secondary Microsoft servers that you select in a failover relationship must be in the same network view. Note the following if the servers are on different network views:
  - You cannot modify or delete failover associations.
  - You cannot add, modify, or delete scopes that are associated with the failover association.
  - The appliance does not display the FQDN, but displays the IP address and status as 'unknown' for the secondary Microsoft server.
  - The appliance may not display the scopes that are assigned to the failover association.
  - The appliance does not allow you to perform any action on the objects in a failover association, as the objects will be in read-only mode.

### Admin Permissions for Managing Microsoft DHCP Failover Relationship

To create, modify or delete a Microsoft failover relationship in NIOS, you must have the same permissions on both the Microsoft servers that are assigned to the failover relationship.

You can update the property of a failover relationship only if it has read/write permissions. You cannot update the properties of the Microsoft server or partner directly. You must first delete the relationship, create a new relationship and assign the DHCP scope to the new relationship. NIOS applies the changes to both Microsoft servers defined in the

failover relationship during its next synchronization schedule, only if they are in read/write mode during that time. You can delete a Microsoft failover relationship only if it is not associated with a DHCP range.

## Microsoft Failover Relationships for DHCP Ranges

When you create or update a DHCP range on NIOS, administrators can assign or remove a DHCP range from a Microsoft failover relationship only if the failover relationship has two read/write Microsoft servers. Note that when you reassign a DHCP range that is assigned to a Microsoft failover relationship, NIOS removes the DHCP range from the failover relationship during the next synchronization of the respective server.

When you remove a DHCP range from a read/write Microsoft failover relationship, NIOS deletes a copy of the DHCP range from both the Microsoft servers defined in the failover relationship. When a DHCP range is associated with a Microsoft failover relationship, any change made to one copy of the range is automatically saved to the other copy.

## Admin Permissions for Managing Failover Relationships for DHCP Ranges

You must have read/write permissions on the DHCP range and on each Microsoft server to assign or remove a DHCP range from a Microsoft failover relationship.

Note that the changes performed on a DHCP range, which is assigned to a read/write Microsoft failover relationship, are applied to both copies of the range and synchronized with each Microsoft server during each of their next scheduled synchronization.

You can delete a DHCP range that is assigned to a Microsoft failover relationship only if the failover relationship has two read/write Microsoft servers. When you delete a DHCP range that is assigned to a Microsoft failover relationship in read/write mode, NIOS deletes them on both Microsoft servers that is defined in the relationship during the next respective scheduled synchronization.

## Limitations of Microsoft DHCP Failover Relationship on NIOS

Note the following limitations of Microsoft DHCP failover relationship on NIOS:

1. You cannot synchronize failover scope statistics on a Microsoft DHCP failover relationship.
2. The appliance does not support on demand replication for a Microsoft DHCP failover relationship.
3. You cannot set the functional peer in the PARTNER-DOWN state for a Microsoft DHCP failover on NIOS. For information about setting a peer in the partner-down state, see [Managing Failover Associations](#).
4. The appliance does not support time synchronization of both partners and time synchronization monitoring. Infoblox recommends that the administrators of the Microsoft servers running the DHCP server must ensure that the time synchronization is appropriate.
5. You can only add up to 31 DHCP failovers to a Microsoft server in a failover relationship.

## Synchronizing IP Addresses with Invalid MAC Addresses

An invalid MAC address is a MAC address whose length is not six bytes (48 bits). You can synchronize IP addresses with invalid MAC addresses only for Microsoft servers for which you have enabled DHCP synchronization. When you enable this setting for a Microsoft DHCP server, NIOS synchronizes any DHCP lease or reservation with an invalid MAC address and it is effective from the next scheduled DHCP lease or DHCP server synchronization. The synchronization does not affect a DHCP lease or reservation with a valid MAC address.

You can synchronize IP addresses with invalid MAC addresses at the Grid level and override them at the member level. This is valid for NIOS 6.11.0 and later releases. The appliance displays the invalid MAC address using a red cross mark in the list. You cannot save invalid MAC addresses, but you can view invalid MAC addresses in the list panels and editors.

When you disable synchronization for invalid MAC addresses, NIOS ignores any DHCP lease or reservation with an invalid MAC address and the object disappears from NIOS during the next synchronization of the Microsoft server to which it belongs. An invalid MAC address that you had synchronized earlier persists until NIOS updates the associated lease or reservation with a new MAC address, or the lease is cleared, or deleted, or expired.

To synchronize an invalid MAC address, complete the following:

1. **Grid:** From the **Grid** tab -> **Grid Manager** tab, expand the Toolbar and click **Grid Properties** -> **Edit**.  
**Member:** From the **Grid** tab -> **Grid Manager** tab, click the **Members** tab, select the *member* checkbox, and click the Edit icon.



2. Select the **Microsoft Server Settings** tab in the *Grid Properties Editor* wizard and complete the following in the **Advanced** tab:
  - **Allow Invalid MAC Address to be synchronized:** This is enabled, by default. Select this to enable synchronization for invalid MAC addresses.  
You can click **Override** at the member level to specify a new value. The **Override** button changes to **Inherit**. Click **Inherit** to inherit the value from the Grid.
3. Save the configuration.

## Managing Synchronized DHCP Data

When Grid members are configured to manage Microsoft DHCP servers in read/write mode, you can use Grid Manager to view, edit and delete the DHCP data of those servers. You can add and manage networks and DHCP ranges that are synchronized as scopes to the Microsoft server, and add and manage reservations and superscopes. All updates are synchronized to the Microsoft servers at regular intervals.

Grid Manager also allows you to set admin permissions, extensible attributes, and thresholds. These apply only when the DHCP data is managed on Grid Manager. These properties are not synchronized to Microsoft servers.

The following sections provide guidelines for managing Microsoft DHCP data from Grid Manager:

- [Adding and Managing Scopes](#)
- [Setting Network Properties](#)
- [Deleting and Restoring a Network](#)
- [Adding a DHCP Range/Scope](#)
- [Setting DHCP Range/Scope Properties](#)
- [Deleting and Restoring a DHCP Range/Scope](#)
- [Viewing Scopes](#)
- [Adding Fixed Addresses/Microsoft Reservations](#)
- [About Superscopes](#)
- [Synchronizing Updates](#)

### Adding and Managing Scopes

To add a scope from Grid Manager, you must create an IPv4 network and a DHCP range, and then assign the Microsoft server to the network and range. To add a split-scope from Grid Manager, you must create an IPv4 network and a DHCP range, and then assign two Microsoft server to the network and range.

To edit a scope synchronized from a Microsoft server, you must edit the properties of its corresponding DHCP range. The following sections describe how to add, edit and remove scopes using Grid Manager.



#### Note

Microsoft servers do not support Infoblox hosts and reservations. You cannot add them to networks and DHCP ranges served by Microsoft servers.

### Adding Networks for Scopes

Following are guidelines for adding a network for Microsoft scopes:

- The network must be served by Microsoft servers. It cannot be served by a mix of Microsoft and Infoblox DHCP servers.
- If you are adding a split-scope, you must assign the network to two Microsoft servers that serve the split-scope. A split-scope cannot be served by a mix of Microsoft and Infoblox DHCP servers.
- The network can contain only one DHCP range per Microsoft server. It can contain multiple DHCP ranges as long as they do not overlap and are each served by a different Microsoft server.
- You can set DHCP properties at the DHCP range level only, not the network level.

You can run discoveries on networks served by Microsoft servers. For information about network discoveries, see [IP Discovery](#) and [vDiscovery](#) below.

 **Note**

Networks served by Microsoft DHCP servers do not support the split, join, and expand functions.

You can create a network from scratch or use a network template. For information about creating network templates, see [Adding IPv4 Network Templates](#). To add an IPv4 network for a scope:

1. From the **Data Management** tab, select the **DHCP** tab.
2. If you have more than one network view in the system, select the network view in which you want to add the network. It must be the same network view to which the Microsoft server is assigned.
3. Expand the Toolbar and click **Add -> Network**.
4. In the *Add Network* wizard, select one of the following and click **Next**:
  - **Add Network**  
or
  - **Add Network using Template**: Click **Select Template** and select a network template. For more information, see [Adding IPv4 Network Templates](#). In the *DHCP Network Template Selector* dialog box, select the template you want to use and click the **Select** icon. Note that when you use a template to create a network, the configurations of the template apply to the new network. The appliance populates the template properties in the wizard when you click **Next**. You can then edit the pre-populated properties, except for **Netmask**.
5. Complete the following and click **Next**:
  - **Address**: Enter the IP address of the network. You can enter the IP address with a CIDR block. For example, enter 10.0.0.0/24, and the netmask slider adjusts the netmask to /24. You can also enter partial IP address with a CIDR block. When you are done, Grid Manager displays the complete IP address with the CIDR block. For example, when you enter 15/24, Grid Manager displays 15.0.0.0/24 and the netmask slider adjusts the netmask to /24. Note that Microsoft DHCP servers do not support /32 subnets.
  - **Netmask**: Use the netmask slider to select the appropriate number of subnet mask bits for the network. Microsoft servers support /1 to /31 netmasks. Note that when you use a template that contains a fixed netmask, you cannot adjust the netmask for this network.
  - **Comment**: Enter additional information about the network, such as the name of the organization it serves.
  - **Automatically create reverse-mapping zone in view**: This function is enabled if the netmask of the network equals /8, /16, or /24. Select this to have the appliance automatically create reverse-mapping zones for the network. A reverse-mapping zone is an area of network space for which one or more name servers have the responsibility for responding to address-to-name queries. These zones are created in the DNS view assigned to receive dynamic DNS updates at the network view level.
  - **Disabled**: This option does not apply to networks assigned to Microsoft servers. The member ignores this field when the network is assigned to Microsoft servers. You can disable DHCP ranges assigned to Microsoft servers, but not networks.
6. Click **Next** to add Microsoft servers as DHCP servers for the network. Click the **Add** icon and select the following:
  - **Add Microsoft Server**: Select the Microsoft server from the *Microsoft Server Selector* dialog box. You can add multiple Microsoft servers, if you are adding multiple DHCP ranges served by different Microsoft servers. For a split-scope, you must assign two Microsoft servers to the network.
7. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [About Extensible Attributes](#).
8. Save the configuration and click **Restart** if it appears at the top of the screen.  
or  
Click the **Schedule** icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Setting Network Properties

You can change the Microsoft servers assigned to the network, and define extensible attributes and admin permissions to the network. You can also set thresholds for the network, to enable the appliance to make a syslog entry when address usage goes above or below the thresholds.

To set network properties:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* checkbox, and then click the **Edit** icon.

2. The *Network* editor contains the following basic tabs from which you can modify data:
  - **General Basic:** You can enter or modify comments.
  - **Member Assignment:** Add or delete Microsoft servers. For information, see as described in [Adding IPv4 Networks](#). If the network contains multiple DHCP ranges each managed by a different Microsoft server, then you can add those Microsoft servers here.
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of the extensible attributes. For information, see [About Extensible Attributes](#).
  - **Permissions:** This tab appears only if you belong to a superuser admin group. For information managing permissions, see [About Administrative Permissions](#).
3. Optionally, you can click **Toggle Expert Mode** to display the following tabs from which you can modify advanced data.
  - **General Advanced:** You can associate zones with a network. For information, see [Associating Networks with Zones](#).
  - **Thresholds:** These watermarks represent thresholds above or below which address usage is unexpected and might warrant your attention. Thresholds are inherited from the Grid.
    - **High-water Mark:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range exceeds this number, the appliance makes a syslog entry. The default is 95.
    - **Low-water Mark:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range drops below this number, the appliance makes a syslog entry. The default is 0. Address usage must initially exceed the low-water mark threshold and then dip below it before the appliance considers low address usage an event requiring an alert.
4. Save the configuration and click **Restart** if it appears at the top of the screen.
 

or

  - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Deleting and Restoring a Network

When you delete a network, Grid Manager moves it and its DHCP ranges and fixed addresses to the Recycle Bin, and permanently deletes its leases. The corresponding scopes and reservations are deleted from the Microsoft server at the next synchronization. If you restore the network on Grid Manager, its DHCP ranges and fixed addresses are restored as well. The Grid member then adds the corresponding scopes and reservations to the Microsoft server on the next synchronization. For information about deleting networks, see [Configuring IPv4 Networks](#). For information about restoring data, see [Using the Recycle Bin](#).

## Adding a DHCP Range/Scope

After you add a network for a scope, you must then define its DHCP range. You can create the DHCP range from scratch or use a DHCP Range template. For information about DHCP templates, see [About IPv4 Range Templates](#).

You can add multiple ranges to the same network, as long as each range is served by a different Microsoft server and the ranges do not overlap.

When you add a split-scope, you must specify the Microsoft servers that serve the scopes and their exclusion ranges. Each scope inherits its options from its respective Microsoft server. Note that the enabled/disabled setting of the first range automatically applies to the second range. Therefore, if the first range is initially disabled, then the second range is initially disabled as well.

To add a DHCP range for a scope:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Navigate to the network to which you want to add a DHCP range, and then click **Add -> DHCP Range** from the Toolbar. You can also add a DHCP range from any panel in the DHCP tab.
3. In the *Add Range* wizard, select one of the following and click **Next**:
  - **Add DHCP Range**
  - or
  - **Add DHCP Range using Template**

Click **Select Template** and select the template that you want to use. Note that when you use a template to create a DHCP range, the configurations of the template apply to the new range. The appliance automatically populates the DHCP range properties in the wizard. You can then edit the pre-populated

- properties.
- Complete the following:
    - Network:** Click **Select Network** to select the network to which you want to add the DHCP range. The network must be served by a Microsoft server. If you are adding a DHCP range while viewing the contents of a specific network, the appliance displays the network address here. You can still select a different network by clicking **Select Network**.
    - Start:** Enter the first IP address in the range.
    - End:** Enter the last IP address in the range.
    - Name:** You can enter a name for the DHCP range.
    - Comment:** You can enter additional information. After the range is synchronized to the Microsoft server as a scope, this text appears in the Description field of the scope on the Microsoft server.
    - Disable for DHCP:** Select this if you do not want the DHCP server to allocate IP addresses from this DHCP range at this time. If you select this, the Grid member synchronizes the range to the Microsoft server as an inactive scope.
  - Click **Next** and select one of the following to provide DHCP services for the DHCP range:

**None:** Select this if you do not want to synchronize this range to the Microsoft DHCP server.

**Microsoft Server:** This field displays the Microsoft server that you selected for the network. If several servers were assigned to the network, you can select one from the list.

- Microsoft Split-Scope:** Select this to create a split-scope, and then complete the following:
    - Microsoft Server #1:** Read-only field that displays the Microsoft server that you specified in the preceding step.
    - Microsoft Server #2:** Select the Microsoft server that will serve the split-scope.
    - Split Percentage:** Specify the percentage of IP addresses in the scope that is allocated to the exclusion range of each Microsoft server. The default is 50%. You can either move the slider or enter the percentages in the text fields. When you use the slider, you are specifying the percentage of addresses in the exclusion range of the first server. A tooltip window displays the percentage as you adjust the slider. When you set the slider, the **Split Percentage**, **Exclusion Starting Address**, and **Exclusion Ending Address** fields are updated accordingly.
    - Exclusion Starting Address:** When you set the split percentages, these fields automatically displays the starting address of the exclusion range of each Microsoft server. Alternatively, you can enter the starting address of the exclusion range of the first Microsoft server, and the **Split Percentage** and **Exclusion Ending Address** values adjust accordingly.
    - Exclusion Ending Range:** When you set the split percentage, these fields automatically display the ending address of the exclusion range of each Microsoft server. Alternatively, you can enter the ending address of the exclusion range of the second Microsoft server, and the **Split Percentage** and **Exclusion Starting Address** values adjust accordingly.
- Click **Next**, and optionally set operational parameters for the scope. Otherwise, the scope inherits its parameters from the first Microsoft DHCP server.
    - Lease Time:** Specify the lease time. The default is 8 days. When the range is served by a Microsoft server and you enter a lease time of 1000 days or more, Grid Manager automatically grays out this field and checks the **Unlimited Lease Time** option after you save your entries.
      - Unlimited Lease Time:** Select this option to set an infinite lease time for the IP addresses leased from this range.
      - Routers:** In the table, enter the IP address of the router that is connected to the same network as the DHCP clients. Click the Add icon to add more routers.
      - Domain Name:** Enter the name of the domain for which the Microsoft server serves DHCP data. The DHCP server includes this domain name in Option 15 when it responds with a DHCPOFFER packet to a DHCPDISCOVER packet from a client.
      - DNS Servers:** In the table, enter the IP address of the DNS server to which the DHCP clients send name resolution requests. The DHCP server includes this information in the DHCPOFFER and DHCPACK messages.
      - Broadcast Address:** Enter the broadcast IP address of the network to which the DHCP server is attached.
  - Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).
  - Save the configuration and click **Restart** if it appears at the top of the screen.  
or

Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Setting DHCP Range/Scope Properties

A Microsoft scope inherits its properties from its Microsoft server. In Grid Manager, you can override the inherited values or set other properties by editing the DHCP range. You can also configure an exclusion range within the scope and set thresholds, to enable the appliance to make a syslog entry when address usage goes above or below the thresholds. You can modify a scope's properties, including its start and end addresses, servers, and exclusion ranges. If you edit the properties of a split-scope and it results in gaps or overlapping exclusion ranges so that the ranges are no longer identical, Grid Manager displays a warning indicating that continuing with the operation automatically removes the split-scope flag. Grid Manager also removes the flag when the start or end address of a scope changes, so its range is no longer the same.

To set DHCP range properties:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr\_range* checkbox, and then click the Edit icon.
2. The *DHCP Range* editor contains the following basic tabs from which you can modify data:
  - **General:** Modify the fields, including the start and end addresses, as described in [Adding a DHCP Range/Scope](#) above.
  - **Server Assignment:** Switch to **None** or select a different Microsoft server for the DHCP range.
  - **IPv4 DHCP Options:** Keep the DHCP properties or override them and enter unique settings for the DHCP range.

This tab displays DHCP and Microsoft vendor options that were synchronized from the Microsoft server. You can edit any of the options. When you select a different User Class or Vendor Class from the drop-down menus, Grid Manager automatically updates the option definitions in the drop-down list. To configure additional DHCP options, click **+** and select a User Class and Vendor Class from the drop-down menus. Select an option from the drop-down list, and enter a value in the field beside it. You can click **-** to remove an option.
  - **Extensible Attributes:** You can add and delete extensible attributes that are associated with a specific DHCP range. You can also modify the values of extensible attributes. For information, see [Managing Extensible Attributes](#).
  - **Permissions:** This tab appears only if you belong to a superuser admin group. For information about managing permissions, see [About Administrative Permissions](#).
3. Optionally, you can click **Toggle Expert Mode** to display the following tabs from which you can modify advanced data.
  - **DDNS:** Complete the following to set DDNS parameters for the range:
    - **Enable DDNS Updates:** Click the checkbox to enable the Microsoft DHCP server to send dynamic DNS updates or clear the checkbox to disable this function.
    - **Option 81 Support**
    - **DHCP Server Updates DNS If Requested by Client:** The DHCP server updates DNS only if it is requested by the client. Otherwise, the client updates DNS.  
**DHCP Server Always Updates DNS:** The DHCP server always updates DNS, regardless of any client request.
  - **Exclusion Ranges:** Configure a range of IP addresses that the server does not use to assign to clients. You can use these exclusion addresses as static IP addresses. For information, see [Configuring IPv4 Fixed Addresses](#). In a split-scope, the exclusion range identifies the range of IP addresses that the other Microsoft server serves. If you edit the exclusion range of either of the scopes in a split-scope and the exclusion ranges no longer complement each other, NIOS removes the split-scope flag from both scopes.
  - **Thresholds:** Thresholds are inherited from the Grid. These watermarks represent thresholds above or below which address usage is unexpected and might warrant your attention.
    - **High-water rMark:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range exceeds this number, the appliance makes a syslog entry. The default is 95.
    - **Low-water Mark:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range drops below this number, the appliance makes a syslog entry. The default is 0. Address usage must initially exceed the low-water mark threshold and then dip below it before the appliance considers low address usage an event requiring an alert.

4. Save the configuration and click **Restart** if it appears at the top of the screen.  
or
  - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Deleting and Restoring a DHCP Range/Scope

When you delete a DHCP range, Grid Manager moves it and its exclusion range and fixed addresses to the Recycle Bin, and permanently deletes its leases. At the next synchronization, the member deletes the scope, its exclusion range and reservations from the Microsoft server. If you restore a DHCP range on Grid Manager, then the Grid member adds its corresponding scope, exclusion range and reservations to the Microsoft server at the next synchronization. For information about deleting DHCP ranges, see [Deleting IPv4 Address Ranges](#). For information about restoring data, see [Using the Recycle Bin](#).

If you delete a scope that is part of a split-scope, Grid Manager automatically removes the split-scope flag from the remaining scope.

## Viewing Scopes

To view the scopes in a network, navigate to **DHCP** -> **Networks** -> *network*. The panel displays the objects in the network, including the scopes and split-scopes. For split-scopes, the panel displays both scopes with the same start and end address. It displays the following information about each object:

- **IP Address:** The IP address of the DHCP object. For a scope, this field displays the start and end addresses of the scope. Note that the appliance highlights all disabled DHCP objects in gray.
- **Split-Scope:** Displays **Yes** if the scope is a split-scope.
- **MS Server:** Displays the Microsoft server that is serving the scope.
- **Type:** The DHCP object type, such as **DHCP Range** or **Fixed Address**.
- **Name:** The object name. For example, if the IP address belongs to a host record, this field displays the hostname.
- **Comment:** The information you entered for the object.
- **IPv4 DHCP Utilization:** The percentage of the total DHCP usage of a DHCP range. This is the percentage of the total number of fixed addresses, reservations, hosts, and active leases in the DHCP range divided by the total IP addresses in the range, excluding the number of addresses in the exclusion ranges. Note that only enabled objects are included in the calculation.
- **Site:** The site to which the DHCP object belongs. This is one of the predefined extensible attributes.

You can select the following additional columns for display:

- **Static Addresses:** Indicates whether the IP address is a static address.
- **Dynamic Addresses:** Indicates whether the IP address is a dynamically assigned address.
- **Disabled:** Indicates whether the object is disabled.
- **Priority:** Displays the priority of a DHCP range when NAC filters are applied.
- Available extensible attributes.

You can also do the following in this panel:

- Sort the displayed data in ascending or descending order by column.
- Click **Go to IPAM View** to view information about the object in the **IPAM** tab.
- Add new objects, such as DHCP ranges, to the network.
- Delete or schedule the deletion of a selected object or multiple objects.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see as described in [Using Quick Filters](#) as described in [Using Quick Filters](#).
- Print or export the data.

You can also view the scopes in the IP Map.



## Adding Fixed Addresses/Microsoft Reservations

To add a reservation from Grid Manager, add a fixed address and Grid Manager synchronizes it to the Microsoft server as a reservation. You can create fixed addresses from scratch or use fixed address templates. For information about fixed address templates, see [Adding IPv4 Fixed Address/Reservation Templates](#).

To add a fixed address:

1. From the **Data Management** tab, select the **DHCP** tab.
2. Expand the Toolbar and click **Add -> Fixed Address**.
3. In the *Add Fixed Address* wizard, select one of the following and click **Next**:
  - **Add Fixed Address**  
or
  - **Add Fixed Address using Template**  
Click **Select Template** and select the template that you want to use.
4. Complete the following:
  - **Network**: Click **Select Network** to select the network to which you want to add the fixed address. If you are adding the fixed address from a specific network, the appliance displays the network address here. You can still select a different network by clicking **Select Network**.
  - **IP Address**: Enter the IPv4 address for the fixed address, or click **Next Available IP** to obtain the next available IP address.
  - **MAC Address**: Enter the MAC address of the host.
  - **Name**: Enter a name for the fixed address. This is required for reservations on Microsoft servers.
  - **Configure On**:
    - None**: Select this if you do not want this synchronized to the Microsoft server.
    - Microsoft Server**: Select the Microsoft server that serves this fixed address.
  - **Comment**: Optionally, enter additional information. The text in this field appears in the Description field of the Microsoft reservation after the fixed address is synchronized. Note that due to a length limit set by the Microsoft DHCP server, after you synchronize DHCP data, the Description field can display only up to 128 characters even though NIOS allows up to 256 characters for this field.
5. Click **Next**, and optionally set operational parameters for the fixed address. Otherwise, the fixed address inherits its parameters from its scope.
  - **Routers**: In the table, enter the IP address of the router that is connected to the same network as the DHCP client. Click the Add icon to add more routers.
  - **Domain Name**: Enter the name of the domain for which the Microsoft DHCP serves DHCP data. The DHCP server includes this domain name in Option 15 when it responds with a DHCP OFFER packet to a DHCPDISCOVER packet from a client.
  - **DNS Servers**: In the table, enter the IP address of the DNS server to which the DHCP client sends name resolution requests. The DHCP server includes this information in the DHCP OFFER and DHCPACK messages.
  - **Broadcast Address**: Enter the broadcast IP address of the network to which the DHCP server is attached.
6. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Using Extensible Attributes](#).
7. Save the configuration and click **Restart** if it appears at the top of the screen.
  - or
  - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Setting Fixed Address/Reservation Properties

Microsoft reservations inherit their properties from their scopes. In Grid Manager, you can override the inherited values or set other properties of a Microsoft reservation, by editing its fixed address.

To modify a fixed address:

1. From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *fixed\_address* checkbox, and then click the Edit icon.
2. The *Fixed Address* editor contains the following basic tabs from which you can enter data:
  - **General**: You can modify the fields described in Adding Fixed Addresses/Microsoft Reservations above.



- **IPv4 DHCP Options:** Keep the inherited properties, or override them and enter unique settings. This section displays DHCP and Microsoft vendor options that were synchronized from the Microsoft server. You can edit any of the options. When you select a different User Class or Vendor Class from the drop-down menus, Grid Manager automatically updates the option definitions in the drop-down list. To configure additional DHCP options, click + and select a User Class and Vendor Class from the drop-down menus. Select an option from the drop-down list, and enter a value in the field beside it. You can click - to remove an option.
- **Discovered Data:** If you ran a discovery on the network, Grid Manager displays the discovered data of the fixed address. For information, see [Viewing Discovered Data](#). Note that conflicts can occur when discovered data does not match the existing IP address data. For information about resolving these conflicts, see [Resolving Conflicting Addresses](#).
- **Extensible Attributes:** Add and delete extensible attributes that are associated with a specific network. You can also modify the values of extensible attributes. For information, see [Managing Extensible Attributes](#).
- **Permissions:** This tab appears only if you belong to a superuser admin group. For information, see [About Administrative Permissions](#).

3. Optionally, you can click **Toggle Expert Mode** to display the **DDNS** tab. To set DDNS parameters for the fixed address, complete the following:

- **Enable DDNS Updates:** Click the checkbox to enable the Microsoft DHCP server to send dynamic DNS updates or clear the checkbox to disable this function.
- **Option 81 Support**
- **DHCP Server Updates DNS If Requested by Client:** The DHCP server updates DNS only if it is requested by the client. Otherwise, the client updates DNS.
- **DHCP Server Always Updates DNS:** The DHCP server always updates DNS, regardless of any client request.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

or

- Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Deleting and Restoring a Fixed Address/Reservation

When you delete a fixed address, Grid Manager moves it to the Recycle Bin. At the next synchronization, the Grid member deletes its corresponding reservation from the Microsoft server. If you restore fixed address, then the Grid member adds its corresponding reservation to the Microsoft server at the next synchronization. For information about deleting fixed addresses, see [Deleting Fixed Addresses](#). For information about restoring data, see [Using the Recycle Bin](#).

## About Superscopes

In Grid Manager, you can group DHCP ranges served by Microsoft servers into a superscope. You can add multiple DHCP ranges to a superscope, as long as the ranges are all served by the same Microsoft DHCP server. The Grid member then synchronizes the superscope and its associated DHCP ranges as superscopes and scopes to the Microsoft DHCP server.

You can also associate extensible attributes with superscopes in Grid Manager. Extensible attributes are not synchronized to the Microsoft DHCP server.

Only admins with read/write permission to superscopes can add and manage superscopes.

## Adding Superscopes

Before you add a superscope, you must first create at least one DHCP range to include in the superscope. To add a superscope:

1. From the **Data Management** tab, select the **DHCP** tab.

2. If you have more than one network view in the system, select the network view in which you want to add the superscope. The network view must be the same one that is assigned to the Microsoft server.
3. Expand the Toolbar and click **Add -> Superscope**.
4. In the *Add Superscope* wizard, complete the following and click **Next**:
  - **Name**: Enter a name for the superscope.
  - **Comment**: Optionally, enter additional information about the superscope.
  - **Disabled**: Select this to disable the DHCP ranges in the superscope. They are then synchronized as inactive scopes on the Microsoft server.
5. Click the Add icon and select a range from the *Select Range* dialog box. This dialog box lists only the address ranges that are served by a Microsoft server.
6. Click **Next** to enter values for required extensible attributes or add optional extensible attributes. For information, see [Managing Extensible Attributes](#).
7. Save the configuration and click **Restart** if it appears at the top of the screen.
 

or

  - Click the Schedule icon at the top of the wizard to schedule this task. In the *Schedule Change* panel, enter a date, time, and time zone. For information, see [Scheduling Tasks](#).

## Viewing Superscopes

To view superscopes, navigate to the **Data Management** tab -> **DHCP** tab -> **Networks** tab -> **Microsoft Superscopes**. Grid Manager displays the following information about each superscope that is displayed:

- **Name**: The name of the superscope. Grid Manager appends the FQDN of its associated Microsoft server so you can identify which superscope belongs to which server.
- **Comment**: The comment that was entered for the superscope.
- **DHCP Utilization**: The percentage of the total DHCP usage of the ranges in the superscope. Fixed addresses and reservations that are outside of a range are excluded from the calculation.
- **Site**: The site of the superscope. This is one of the predefined extensible attributes.

You can add the following columns for viewing:

- **Static Addresses**: The number of static addresses.
- **Dynamic Addresses**: The number of dynamic addresses.
- **Disabled**: Indicates whether the superscope is enabled.

You can do the following in this section:

- Click the link of a superscope to list its address ranges.
- Add a superscope.
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information about using quick filters, see [Finding and Restoring Data](#).
- Print or export the information in this section.
- Delete a superscope.

## Modifying Superscopes

To modify a superscope:

1. From the **Data Management** tab, select the **DHCP** tab -> **Network** tab -> **Microsoft Superscopes** -> *ms\_superscope* checkbox, and then click the Edit icon.
2. The *Superscopes* editor contains the following tabs from which you can modify data:
  - **General**: You can modify the name and comment, and enable or disable the superscope. You can also add and delete address ranges from the superscope. Note that when you delete the last DHCP range in a superscope, Grid Manager automatically deletes the superscope as well.

- **Extensible Attributes:** Define extensible attributes for the superscope. These apply only when the superscope is managed in Grid Manager. For information, see [Managing Extensible Attributes](#).
  - **Permissions:** Define administrative permissions that apply to the superscope when it is managed in Grid Manager. For information see [About Administrative Permissions](#).
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Deleting Superscopes

When you delete a superscope in Grid Manager, it is permanently deleted from the database. The superscope is deleted from the Microsoft server at the next synchronization. Note that deleting a superscope does not delete the DHCP ranges in the superscope. These are retained in the database.

To delete a superscope:

1. From the **Data Management** tab, select the **DHCP** tab -> **Network** tab -> **Microsoft Superscopes** -> *ms\_superscope* checkbox, and then click the Delete icon.
2. Click **Yes** when the confirmation dialog appears.

## Synchronizing Updates

A Grid member synchronizes DHCP data with each of its managed Microsoft server at regular intervals. During each synchronization, updates from Grid Manager are applied to the Microsoft server and updates from the Microsoft server are applied to the Grid as well.

Because admins can update DHCP data from both the Microsoft server and from Grid Manager, conflicts can occur during synchronization. The following guidelines describe how the Grid member resolves conflicts and handles any differences when DHCP data is synchronized between a Microsoft server and the Grid.

- If a Microsoft server admin modifies an object that has a pending scheduled task in Grid Manager and synchronization occurs before the scheduled task, the object is modified in both the Microsoft server and the Grid member. When the scheduled task executes at its scheduled time, it fails and an error message is logged in the audit log.
- When a Microsoft server admin and a Grid Manager admin change the same object, the Grid member retains the version that exists on the Microsoft server. Following are some examples:

Grid Manager Admin...	Microsoft Server Admin...	After Synchronization
Deletes the 10.1.1.0/24 network which has two DHCP ranges	Adds a scope that is within the 10.1.1.0/24 network	The 10.1.1.0/24 network is created on the Grid with the updates and is assigned to the Microsoft server.
Changes the DHCP options of a scope	Deletes the scope.	The scope is deleted from the Grid as well.

- If a Grid member manages multiple Microsoft servers, it can synchronize scopes to the same network as long as they are served by different Microsoft servers and they do not overlap. If the Microsoft servers have scopes that overlap, the Grid member synchronizes only one of the scopes, including its reservations. It does not synchronize the other scopes and logs an error message for each scope that is not synchronized. For information about the Microsoft logs, see [Viewing Synchronization Logs](#).  
Note that a Grid member can synchronize scopes with overlapping reservations because they are served by different Microsoft servers.
- When a Grid member synchronizes a split-scope to its respective Microsoft servers, the scopes use the default value for the DHCP Offer Delay value, since this property is not supported by NIOS.
- If you create a split-scope on a NIOS appliance, synchronization fails if there is an existing scope in the same network on one of the Microsoft servers. Only one scope is allowed in a network, per Microsoft server.
- If a Microsoft admin adds a DHCP range and a NIOS admin is in the process of adding the same range when a synchronization occurs, the NIOS admin will not be able to save the range after the synchronization. Grid Manager will display an error message indicating that the range already exists.
- If both a NIOS admin and a Microsoft admin create a scope or split-scope and conflicts occur, the Microsoft server always takes precedence. All conflicts are logged to the Microsoft log. Following are some examples:

- If the NIOS admin creates a scope and a Microsoft server admin creates a split-scope for the same DHCP range, the split-scope is synchronized to Grid Manager.
- If the NIOS admin creates a split-scope on Microsoft servers 1 and 2, and a Microsoft admin creates the same split-scope on Microsoft servers 1 and 3 but with different exclusion ranges, the scope created by the NIOS admin on Microsoft server 1 is dropped upon synchronization.
- If the NIOS admin creates a split-scope on Microsoft servers 1 and 2, and a Microsoft admin creates the same split-scope on the same Microsoft servers but with different exclusion ranges, the split-scope created by the Microsoft admin is synchronized to NIOS and retained. The split-scope created by the NIOS admin is dropped.

## Synchronizing DHCP Data from Microsoft Servers

After you configure a member to manage the DHCP service of a Microsoft server, the Grid member connects to the server and starts synchronizing IPv4 DHCP data from the Microsoft server to the Grid database. It synchronizes the Microsoft server properties, leases, scopes and reservations.

The synchronization time varies, depending on different factors, such as the number of managed Microsoft servers and the amount of data being synchronized.



### Note

Synchronizing IPv6 data is not supported.

As shown in the table below, Microsoft servers and Infoblox DHCP servers represent DHCP data differently. Scopes on Microsoft servers are DHCP ranges on Infoblox DHCP servers. Additionally, Microsoft servers support split-scopes, which is a scope assigned to two Microsoft servers. Each scope has an exclusion range on opposite ends to specify the pool of IP addresses that the other Microsoft server allocates. On an Infoblox DHCP server, each scope in the split-scope is represented as a DHCP range with an exclusion range. Note that NIOS also synchronizes scopes assigned to more than two Microsoft servers, but they are not synchronized as split-scopes.

Fixed addresses on Infoblox DHCP servers are the same as reservations on Microsoft servers. Infoblox reservations, which are IP addresses that are excluded from DHCP, are not supported on Microsoft servers. Microsoft superscopes, which are used to group scopes, are represented as superscopes and can be managed from Infoblox DHCP servers.

### *DHCP Data in Microsoft and Infoblox DHCP Servers*

DHCP Data	Microsoft DHCP Servers	Infoblox DHCP Servers
Address pool from which the server allocates addresses	Scope	DHCP Address Range in a Network
An IP address that is always assigned to the same device	Reservation	Fixed Address
An IP address that is excluded from DHCP because a user intends to configure it manually on a network device	Not supported	Reservation
Administrative group of scopes	Superscope	Microsoft superscope



### Note

In this chapter, reservations always refer to Microsoft reservations (Infoblox fixed addresses), unless otherwise specified.

When the member synchronizes a scope to the Grid, it converts the scope to a DHCP range and network. For example, it converts the Microsoft scope 10.1.1.1- 10.1.1.200 with a netmask of /24 to the network 10.1.1.0/24 and DHCP range 10.1.1.1- 10.1.1.200 on Grid Manager. The member associates the DHCP properties of the scope, including its DHCP and Microsoft vendor options, with the DHCP range. It synchronizes the leases within the range and if configured, the

exclusion range as well.

NIOS synchronizes two scopes as split-scopes if the following conditions are met:

- Two scopes have the same address range.
- The scopes are assigned to two different Microsoft servers.
- Each scope has an exclusion range and the exclusion ranges are at opposite ends of the scope, so they complement each other. For example, the scope 10.1.1.1-10.1.1.200 on Microsoft server A has an exclusion range of 10.1.1.100-10.1.1.200 and the same scope on Microsoft server B has an exclusion range of 10.1.1.1-10.1.1.99.

When the appliance synchronizes a split-scope, it sets a split-scope flag on each scope to indicate that it is part of a split-scope. For more information, see [Viewing Scopes](#). It synchronizes any reservations that are configured in each scope as well.

When the member synchronizes a Microsoft reservation to the Grid, it converts the reservation to a fixed address and static lease on Grid Manager. It associates the DHCP properties and DHCP and Microsoft vendor options of the reservation with the fixed address record.

The Grid member synchronizes superscopes to the Grid as well. The Grid supports Microsoft superscopes, when an MS management license is installed. For information about adding and managing superscopes in Infoblox DHCP servers, see [About Superscopes](#).

Following are some guidelines on how a Grid member synchronizes DHCP data from Microsoft servers to the Grid:

- If two superscopes have the same name, but are served by different servers, the member creates two different superscopes on the Grid, each appended with the Microsoft server FQDN.
- The member synchronizes all active and inactive scopes from a managed Microsoft server as long as the scopes do not conflict or include any networks currently served by a Grid member. The member does not synchronize a scope if its network already exists in the Grid and is served by a Grid member. It can synchronize a scope if its network is included in an existing network, only if the network is not served by DHCP.
- Synchronizing scopes that are larger than /12 is not supported.
- NIOS synchronizes all scopes except for those with serving ranges that overlap the serving ranges of existing DHCP ranges.
- If the appliance manages multiple Microsoft servers and synchronizes identical scopes from more than two Microsoft servers, it does not flag the scopes as split-scopes.
- If the appliance synchronizes one or more scopes from Microsoft servers that are identical to an existing split-scope, it removes the split-scope flag from the existing split-scope.
- NIOS does not synchronize partially overlapping scopes inside a single network from different Microsoft servers. It synchronizes only ranges that completely overlap.
- More than two scopes are not synchronized as split-scopes, even if they are identical and have exclusion ranges that complement each other.
- Scopes that have more than one exclusion range are not synchronized as split-scopes, even if the exclusion ranges complement each other. In addition, if a split-scope is synchronized from a Microsoft server and one of the scopes is split again on the Microsoft server, NIOS synchronizes the third scope, but does not set a split-scope flag. In addition, it removes the split-scope flag from the original split-scopes.

You can view the synchronized data as follows:

- To view the networks of the scopes, select the **Data Management** tab -> **DHCP** tab -> **Networks** tab -> **Networks** panel. This panel displays all IPv4 networks.
- To view the corresponding DHCP ranges and reservations, select the **Data Management** tab -> **DHCP** tab -> **Networks** tab, and click a network link. For information about this panel, see [Viewing Scopes](#).

You can also use the features in the **IPAM** tab, such as the Net Map and IP Map, to view and manage the Microsoft DHCP data. For information, see [About IP Address Management](#).

## Managing Microsoft DHCP Servers

You can control the DHCP services of managed Microsoft servers and set certain properties as well. This section includes the following topics:

- [Setting Microsoft DHCP Server Properties](#)
- [Controlling the DHCP Service of a Microsoft Server](#)

- [Disabling and Removing Microsoft DHCP Servers](#)
- [Modifying DHCP Server Assignments](#)

## Viewing Members and Managed DHCP Servers

You can view Infoblox and Microsoft DHCP servers by navigating to the **Data Management** tab -> **DHCP** tab, and then selecting the **Members/Servers** tab. The panel displays the following information about each DHCP server:

- **Name:** The hostname of the Grid member or Microsoft server.
- **Status:** The status of the DHCP service on the Grid member or Microsoft server.
- **Comment:** Comments that were entered for the Grid member or Microsoft server.
- **DHCP Utilization:** The percentage of the total DHCP utilization of the member or Microsoft server. This is the percentage of the total number of DHCP hosts, fixed addresses, reservations, and leases assigned to the member or Microsoft server versus the total number of IP addresses (excluding IP addresses in the exclusion range) and all DHCP objects assigned to the member or DHCP server. Note that only enabled objects are included in the calculation. The appliance updates the utilization data every 15 minutes. The appliance displays the utilization data in one of the following colors:
  - Red: The DHCP resources are 100% utilized.
  - Yellow: The utilization percentage is over the effective high watermark threshold.
  - Blue: The utilization percentage is below the effective low watermark threshold.
  - Black: The utilization percentage is at any number other than 100%, or within the effective thresholds.
- **Site:** Values that were entered for this pre-defined attribute.

You can select the following additional columns for display:

- **Address:** The IP address of the member or Microsoft server.
- **Static Addresses:** The number of static IP addresses.
- **Dynamic Addresses:** The number of dynamically assigned IP addresses.

You can do the following:

- Use filters and the **Goto** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Goto** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Edit the properties of a Grid member or Microsoft server.
  - Click the checkbox beside a Grid member or Microsoft server, and then click the Edit icon.
- Export the list of Grid members and Microsoft servers to a .csv file.
  - Click the Export icon.
- Print the list of Grid members and Microsoft servers.
  - Click the Print icon.

## Setting Microsoft DHCP Server Properties

From Grid Manager, you can set DHCP properties supported by a Microsoft server. These are applied to the server at the next synchronization. You can also set other properties that apply to Grid Manager only, such as thresholds and the logging.

To set properties for a Microsoft DHCP server:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members/Servers** tab -> **Members/Servers** -> *ms\_server* checkbox, and then click the Edit icon.
2. In the *Microsoft Server DHCP Properties* editor, you can configure DHCP properties in each tab as follows: **IPv4 DHCP Options** tab: Complete the following to configure basic DHCP options for the server:
  - **Routers:** Click the Add icon and enter the IP address of the router that is connected to the same network as the DHCP clients.
  - **Domain Name:** Enter the name of the domain for which the server serves DHCP data. The DHCP server includes this domain name in Option 15 when it responds with a DHCP OFFER packet to a DHCPDISCOVER packet from a client. If DDNS is enabled on the DHCP server, it combines the host name from the client and this domain name to create the FQDN (fully-qualified domain name) that it uses to update DNS.



- **DNS Servers:** Click the Add icon and enter the IP address of the DNS server to which the DHCP client sends name resolution requests. The DHCP server includes this information in the DHCP OFFER and DHCP ACK messages.
- **Broadcast Address:** Enter the broadcast IP address of the network to which the DHCP server is attached.
- **Custom DHCP Options:** This section displays DHCP and Microsoft vendor options that were synchronized from the Microsoft server. You can edit any of the options. When you select a different User Class or Vendor Class from the drop-down menus, Grid Manager automatically updates the option definitions in the drop-down list.  
To configure additional DHCP options, click + and select a User Class and Vendor Class from the drop-down menus. Select an option from the drop-down list, and enter a value in the field beside it. You can click - to remove an option.

**DDNS** tab: You can enable or disable dynamic DNS updates and set certain properties.

- **Enable DDNS Updates:** Click the checkbox to enable the Microsoft DHCP server to send dynamic DNS updates or clear the checkbox to disable this function.
- **Option 81 Support**  
**DHCP Server Updates DNS If Requested by Client:** The DHCP server updates DNS only if it is requested by the client. Otherwise, the client updates DNS.  
**DHCP Server Always Updates DNS:** The DHCP server always updates DNS, regardless of any client request.

**Thresholds** tab: Thresholds are inherited from the Grid. These watermarks represent thresholds above or below which address usage is unexpected and might warrant your attention.

- **Enable DHCP Thresholds:** Select this checkbox to enable the feature.
    - **High-water Mark:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range exceeds this number, the DHCP server makes a syslog entry. The default is 95.
    - **Low-water Mark:** Enter a number between 0 and 100. If the percentage of allocated addresses in a DHCP range drops below this number, the DHCP server makes a syslog entry. The default is 0. Address usage must initially exceed the low-water mark threshold and then dip below it before the appliance considers low address usage an event requiring an alert.
3. Optionally, you can click **Toggle Expert Mode** to display the **Logging** tab, where you can enable the managing member to log the lease events of the Microsoft server. This setting is inherited from the Grid. You can override that setting by clicking **Override**, and then selecting or clearing the **Log Lease Events from DHCP server** checkbox.
  4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Controlling the DHCP Service of a Microsoft Server

You can start and stop the DHCP service of a managed Microsoft server from Grid Manager as follows:

1. From the **Data Management** tab, select the **DHCP** tab -> **Members/Servers** tab -> **Members/Servers** -> *ms\_server* checkbox.
2. Expand the Toolbar and click **Start** or **Stop**.
3. Click **Yes** when the confirmation dialog appears.

## Disabling and Removing Microsoft DHCP Servers

If you remove a Microsoft server as a managed server, Grid Manager deletes all the DHCP ranges, leases, and fixed addresses associated with the server. It also deletes networks that were assigned only to the Microsoft server. It does not delete a network if it was assigned to other Microsoft servers as well.

When you disable a Microsoft server, the managing Grid member terminates any on-going synchronization and restarts synchronization only when the server is re-enabled. The DHCP data associated with that server is preserved in the same state until synchronization resumes.

For information on removing and disabling Microsoft servers, see [Managing Microsoft Servers](#).



## Modifying DHCP Server Assignments

If you disable a Microsoft DHCP server or take it offline for maintenance purposes, for example, you can assign its scopes to a member DHCP server.

Following are the tasks to reassign scopes from a Microsoft server to a member DHCP server:

1. Set the server assignments of all fixed addresses in the scope to "None".  
From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *fixed\_address* checkbox, and then click the Edit icon. You can change the server assignment in the **General** tab of the *Fixed Address* editor.
2. Set the server assignments of all address ranges served by the Microsoft server to "None".  
From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr\_range* checkbox, and then click the Edit icon. You can change the server assignment in the **General** tab of the *DHCP Range* editor.
3. Change the sever assignments of the networks by deleting the Microsoft server and replacing it with a member DHCP server.  
From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* checkbox, and then click the Edit icon. You can change the server assignment in the **Member Assignment** tab of the *Network* editor. contains the following basic tabs from which you can modify data:
4. Modify the server assignments of all address ranges and specify the member DHCP server.  
From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr\_range* checkbox, and then click the Edit icon. You can change the server assignment in the **General** tab of the *DHCP Range* editor.
5. Restart services.

The member DHCP server starts granting lease requests after the restart. Note that you do not need to clear the leases that were active on the Microsoft server, because the member automatically clears them when you change the DHCP server assignment.

## Monitoring

This section explains how to use the different monitoring and reporting tools, including DHCP fingerprint detection and SNMP. It includes the following topics:

- [Monitoring the Appliance](#)
- [Configuring DHCP Fingerprints and Fingerprint Detection](#)
- [Monitoring with SNMP](#)

## Monitoring the Appliance

Grid Manager provides tools for monitoring the status of the Grid, members, and services. You can monitor overall Grid and member status from the Dashboard, which provides a high-level view of your Grid, members and IP address data, and easy access to tasks. For information, see [Dashboards](#).

Grid Manager also displays status icons to indicate the state of appliances, services, database capacity, Ethernet ports, HA, and Grid replication. Depending on your appliance, Grid Manager can display status icons for the power supplies as well as icons to indicate the state of the RAID array and disk controller backup battery.

You can monitor detailed status of the Grid, members, and services, and then decide how to manage them. Note that when any member or service encounters issues, the appliance sends SNMP traps. For information, see [Monitoring with SNMP](#).

This section describes the status icons that indicate the state of appliances, services, database capacity, Ethernet ports, HA, and Grid replication. It also explains how to use the various logs and the traffic capture tool to monitor a NIOS appliance. It contains the following topics:

- [Viewing the Grid Node Tree](#)

- [Member Status](#)
- [Enabling Grid Visualization](#)
- [Monitoring Services](#)
- [Capturing DNS Queries and Responses](#)
- [Monitoring Tools](#)
- [Using a Syslog Server](#)
- [Enabling Automated Traffic Capture](#)
- [Tracking Object Changes in the Database](#)
- [Configuring the PTop Instance to Collect CPU Utilization Data](#)
- [Collecting Database Performance Data](#)
- [Configuring dnstap](#)

## Viewing the Grid Node Tree

You can view graphical representation of the Grid, with its members represented as nodes in the tree. The **Visualization** tab is disabled for new installations and this tab is displayed by default when you upgrade to NIOS 7.2.x and later. For information about enabling this option, see [Enabling Grid Visualization](#). Each member is labeled with its hostname. You can click **Display Node Labels** on the left panel to display or hide the labels.

By default, the Grid Master is the root node at the center of the tree. It is represented by a color-coded icon connected to its members. You can then click a member to re-center the tree on that node. The left panel displays information about the member that is at the center of the node tree.

In the node tree, the shape of the icons indicate the role of the member in the Grid:

- Circle: Grid Master
- Ellipse: Grid Members

The colors of the icons indicate the status of the member:

- Green: The member is online and functioning properly.
- Grey: The member has not joined the Grid.
- Red: The member has operational problems.

The connectors indicate the connection status between the Grid Master and the member.

- Blue Line: Connects the Grid Master with online Grid members
- Thick White Line: Connects the Grid Master with Grid Master Candidates
- Dashed Line Connector: Connects the Grid Master with offline Grid members

The node tree includes zooming and panning capabilities to enable quick navigation and selection among multiple nodes. You can also hover your mouse over a node to view node information. It displays the same information as that displayed on the left panel, when a node is at the center of the tree.

For the Grid Master:

- Grid name
- Standalone or HA
- Number of members in the Grid
- Status of each protocol running on the Grid
- Grid status For a Member:
- Member name
- Standalone or HA
- HA Status if HA pair
- Status of each protocol running on that member

## Member Status

You can monitor the overall status, such as the memory usage and system temperature, of a Grid member or an independent appliance using the *Member Status (System Status)* widget on the Dashboard.

To monitor the detailed status of a member, from the **Grid** tab, select the **Grid Manager** tab -> **Members** tab.

In the **Members** tab, Grid Manager displays the Grid Master first and then all other members in alphabetical order. If a member is an HA pair, you can click the arrow next to the member row to view information about the active and passive nodes. Grid Manager can display the following information:

- **Name:** The name of the member.
- **HA:** Indicates whether the member is an HA pair.
- **Status:** The service status of the member. For a vNIOS appliance whose license is revoked and is still operating in the Grid, Grid Manager displays a license violation warning here. You should immediately remove this member from the Grid.
- **IPv4 Address:** The IPv4 address of the appliance, or the VIP of an HA pair.
- **IPv6 Address:** The IPv6 address of the appliance, or the VIP of an HA pair.
- **Identify:** This field appears only if your appliance has the unit identification button. This can be On or Off. When you identify the appliance by pressing the UID button on the appliance or through the GUI or CLI command, this field displays On. Otherwise, this is Off.
- **DHCP, DNS, TFTP, HTTP, FTP, NTP, bloxTools, Captive Portal, DNS Accelerator usage, Reporting, Discovery, Threat Protection, Cloud-API, Threat Analytics, TAXII:** The status icons indicate whether these services are running properly. For information about service status, see [Monitoring Services](#).
- **Hardware Type:** The hardware type of the appliance.
- **Hardware Model:** The hardware model of the appliance.
- **Serial Number:** The serial number of the appliance.
- **DB Utilization:** The current percentage of the database in use.
- **Host Platform:** The platform on which the appliance is running. For a vNIOS appliance, this field displays the name of the cloud management platform, such as **AWS, Azure, GCP, or VMware**, and for the vNIOS for GCP appliance, if the NIOS instance is running on a single network interface, the field displays **GCP (Single Interface)**. For physical NIOS appliances, this displays **Physical**. For appliances running on the CentOS operating system, the **Host Platform** column displays **Red Hat**.
- **Hypervisor:** The hypervisor of the appliance
- **Comment:** Information about the member.



**Note**

The placement of the **Host Platform** and **Hypervisor** columns may vary after a NIOS upgrade.


To turn the identification button on or off on the member, click the Hardware Identify icon from the horizontal navigation bar. Grid Manager displays a panel with the appliance name, status, and IP address. Hover your mouse over the row and click **Turn On** to turn the identification button on, or click **Turn Off** to turn it off.



To view detailed status, select a member checkbox, and then click the Detailed Status icon. Grid Manager displays the *Detailed Status* panel. If the selected member is an HA pair, Grid Manager displays the information in two columns, one for the active node and the other for the passive. The *Detailed Status* panel provides detailed information described in the following sections.

You can modify some of the data in the table. Double click a row, and either modify the data in the field or select an item from a drop-down list. Click **Save** to save the changes. Note that some fields are read only.

**Appliance Status**

The status icon indicates the operational status of a Grid member and a general description of its current operation. The status icon can be one of the following:




Icon	Color	Meaning
	Green	The appliance is operating normally in a "Running" state.

	Yellow	The appliance is connecting or synchronizing with its Grid Master.
	Red	The Grid member is offline, is not licensed (that is, it does not have a DNSone license with the Grid upgrade that permits Grid membership), is upgrading or downgrading, or is shutting down.

The following are descriptions that may appear: Running, Offline, Error, and Warning.



### Disk Usage

Grid Manager displays the percentage of the data partition of the hard disk drive that is currently in use on the selected Grid member. It also displays whether the percentage of usage has exceeded the trigger or reset value. Note that the trigger and reset values are user configurable. The default trigger value is 85% and reset value is 70%. The status icon can be one of the following:

Icon	Color	Meaning
	Green	The disk usage is either below the reset value or has not yet reached the trigger value.
	Yellow	The disk usage is decreasing from the trigger value, but has not yet reached the reset value.
	Red	The disk usage has exceeded the trigger value.




### DB Capacity Usage

Grid Manager displays the current percentage of the database in use on the selected Grid member. It also describes whether the usage has exceeded the trigger or reset value. Note that the trigger and reset values are user configurable. The default trigger value is 80% and the reset value is 70%. For information about using the capacity report, see [Monitoring Tools](#). The status icon can be one of the following:

Icon	Color	Meaning
	Green	The database capacity is either below the reset value or has not yet reached the trigger value.
	Yellow	The database capacity is decreasing from the trigger value, but has not yet reached the reset value. When the capacity exceeds the trigger value, the icon changes from green to yellow.



### LAN1/LAN2 Ports, HA Port, and MGMT Port

Grid Manager displays the IP address of the port. The status icons for these ports indicate the state of their network connectivity.

Icon	Color	Meaning
	Green	The port is properly connected to a network. Grid Manager displays the IP address of the network.
	Red	The port is not able to make a network connection.
	Gray	The port is disabled.




## LCD

The LCD status icon indicates its operational status.

Icon	Color	Meaning
	Green	The LCD is functioning properly.
	Yellow	The LCD process is not running.




## Memory Usage

Grid Manager displays the current percentage of system memory in use on the selected Grid member. It also describes whether the usage has exceeded the trigger or reset value. Note that the trigger and reset values are user configurable. The default trigger value is 90% and the reset value is 80%. You can see more details about memory usage through the CLI command: `show memory`. The status icon can be one of the following.

Icon	Color	Meaning
	Red	The memory usage has exceeded the trigger value.
	Yellow	The memory usage is decreasing from the trigger value, but has not yet reached the reset value.
	Green	The memory usage is either below the reset value or has not yet reached the trigger value.

## Swap Usage

Grid Manager displays the current percentage of swap area in use on the selected Grid member. It also describes whether the usage has exceeded the trigger or reset value. Note that the trigger and reset values are user configurable. The default trigger value is 20% and the reset value is 10%. The status icon can be one of the following:

Icon	Color	Meaning
	Red	The memory usage has exceeded the trigger value.
	Yellow	The memory usage is decreasing from the trigger value, but has not yet reached the reset value.
	Green	The memory usage is either below the reset value or has not yet reached the trigger value.



## FAN

The status icon indicates whether the fan is functioning properly. The corresponding description displays the fan speed. The status icon and fan speed are displayed for Fan1, Fan2, and Fan3.





### Note



vNIOS appliances on VMware do not monitor or report the fan speed.

Icon	Color	Meaning
	Green	The fan is functioning properly.
	Red	The fan is not running.

## NTP Synchronization




The status icon indicates the operational status of the current NTP synchronization status.

Icon	Color	Meaning
	Green	The NTP service is enabled and running properly.
	Yellow	The NTP service is enabled, and the appliance is synchronizing its time.

	Red	The NTP service is enabled, but it is not running properly or is out of synchronization.
	Gray	The NTP service is disabled.

### Passive HA Connectivity Status

The status icon indicates the ARP connectivity status of the passive node of an HA pair. The status icon can be one of the following:

Icon	Color	Meaning
	Green	The passive HA node is connected to the local router.
	Yellow	The passive HA node fails to connect to the local router.
	Gray	ARP is disabled on the passive node of an HA pair.

### CPU Temperature

This icon is always green. The description reports the CPU temperature.



#### Note

vNIOS appliances on VMware do not monitor or report the CPU temperature.

### System Temperature

This icon is always green. The description reports the system temperature.



#### Note

vNIOS appliances on VMware do not monitor or report the system temperature.



### CPU Usage

Grid Manager displays the current percentage of the CPU usage on the selected Grid member. The maximum is 100%. It also describes whether the CPU usage has exceeded the trigger or reset value. Note that the trigger and reset values are user configurable. The default trigger value is 81% and the reset value is 70%. You can see more details about CPU usage through the CLI command: show CPU.

The status icon can be one of the following:

Icon	Color	Meaning
------	-------	---------



	Green	The CPU usage is either below the reset value or has not yet reached the trigger value.
	Red	The CPU usage has exceeded the trigger value.

## Enabling Grid Visualization

To enable Grid visualization:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Grid Properties** -> **Edit** from the Toolbar.
2. In the *GridProperties* editor, click the **Advanced** tab.
3. Select the **Show Grid Visualization** checkbox.
4. Click **Save** to save the changes. The **Visualization** tab appears and displays a graphical representation of the Grid. This may take a long time depending on the number of members you have in the Grid. When you deselect this checkbox, you disable Grid visualization for all users.

## Grid Status

You can monitor the overall status of the Grid using the *Grid Status* widget on the Dashboard.

You can also view the Grid status from the **Grid Manager** tab. To view Grid status, from the **Grid** tab, select the **Grid Manager** tab. Grid Manger displays the overall Grid status and status of all Grid services. The Grid status represents the status of the most critical members or services in the Grid. When all Grid members are running properly, the overall Grid status is green. When one of the members has operational problems, the overall Grid status is red. Grid Manager lists all Grid members in the **Members** tab so you can identify which member has issues. For information, see [Member Status](#).

In addition, the service bar below the Grid status lists the status of all licensed services on your Grid. This can include DHCP, DNS, Cloud-API, TFTP, HTTP (File Distribution), FTP, NTP, bloxTools, Captive Portal, DNS Accelerator, Reporting, Discovery and others, depending on the active licenses you have installed. When you click a service link, Grid Manager displays detailed information about the selected service running on all members. For information, see [Monitoring Services](#).


Grid Manager also provides icons you can use to edit Grid properties and bookmark the page.




## Monitoring Services

The Grid or device status icon and the service icon indicates whether a service running on a member or an independent appliance is functioning properly or not.

### Service Status

After you enable any of the services — DHCP, DNS, TFTP, HTTP (for file distribution), FTP, NTP, bloxTools, Captive Portal, Reporting, Discovery, Threat Protection and Cloud API — the appliance indicates their status as follows:

Icon	Color	Meaning
	Green	The service is enabled and running properly.

	Yellow	The service is enabled, but there may be some issues that require attention. For Threat Protection, this could mean that one of the members is in monitor mode.
	Red	The service is enabled, but it is not running properly. (A red status icon can also appear temporarily when a service is enabled and begins running, but the monitoring mechanism has not yet notified Grid Manager.)
	Gray	The service is not configured or it is disabled.

 **Note**

When you enable reporting service on the Grid and configure multi-site cluster, you can monitor the status of all reporting members that you have configured. For information about reporting clusters, see [Configuring Reporting Clustering](#).

## Monitoring Grid Services

The status icon of a Grid service represents the status of the most critical service in the Grid. For example, if the Grid DHCP status icon is red, the DHCP service on one of the members in the Grid is not running properly. You can click the DHCP service link to view the service status of all Grid members and identify which member has a service problem. You can then decide to start or stop the service, or modify the service configuration on that member.

To monitor a Grid service:

1. From the **Grid** tab, select the **Grid Manager** tab, and then click a service link.
2. Grid Manager displays the following information in the **Services** tab:
  - **Name:** The name of the member.
  - **Service Status:** The current status of the service.
  - **IPv4 Address:** The IPv4 address of the appliance or the VIP of an HA pair.
  - **IPv6 Address:** The IPv6 address of the appliance or the VIP of an HA pair.
  - **Comments:** Information about the member or service.
  - **Site:** The site to which the member belongs. This is one of the predefined extensible attributes. You can select available extensible attributes for display.
  - **Reporting Site:** This field appears only when you enable the reporting service and configure the multi-site cluster. For information about how to configure the multi-site clustering mode, see [Configuring Reporting Clustering](#).

 **Note**

The **Reporting Site** column is hidden by default. To display this column, click the down arrow next to any column header and select **Columns** -> **Edit Columns** -> **Reporting Site** checkbox and click **Apply**. If the **Reporting Site** column is visible, then the extensible attribute value is automatically updated.

3. Optionally, click the Edit icon next to the service name to edit the Grid properties for the service.

or

Select a member checkbox, and do one of the following:

- Click the Edit icon to edit the member service configuration. Grid Manager displays the editor for the corresponding service. For example, when you edit the DHCP service, Grid Manager displays the *Member DHCP Configuration* editor.
- Click the Start icon to start the service.
- Click the Stop icon to stop the service.
- Grid Manager updates the service status based on your action.

## Monitoring Member Services

You can view detailed service status on a selected member. Optionally, you can start and stop a service, and edit the service configuration.

To monitor a member service:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox, and then click the Manage Member Services icon.  
In the *Manage Services* panel, Grid Manager displays the following information:
  - **Service:** The name of the service.
  - **Status:** The current status of the service running on the member.
  - **Description:** The description of the status. Grid Manager displays the percentage of usage for the TFTP, HTTP (File Distribution), FTP, and bloxTools services.
2. Optionally, mouse over a service and do one of the following:
  - **Start/Stop Service:** Click this icon to start or stop the selected service. For example, when the DNS service is currently stopped, the appliance starts the service when you click this icon.
  - **Edit Service:** Click this icon to edit the selected service. Grid Manager displays the corresponding editor. For example, when you click the Edit Service icon for DNS, Grid Manager displays the *Member DNS Configuration* editor.
3. Click the Refresh icon to update the service status.

## Capturing DNS Queries and Responses

You can capture DNS queries and responses for later analysis. When configuring this feature, you can choose to save the capture file locally on your appliance, as well as on the FTP (File Transfer Protocol) or SCP (Secure Copy) server. When you save it locally, you can use `show query_capture` to view the contents of the capture file. You can also use filter commands to exclude certain queries and view only the desired ones. Note that using multiple CLI commands to filter data for the appliances with large number of captured DNS queries and responses can significantly affect the system performance, protocol performance, and CLI command performance.

A capture file for logging DNS queries and responses is rolled over based on the configured time limit or when the file reaches 100 MB in size, whichever is sooner. The default time limit is 10 minutes. The capture file is automatically saved and exported to an FTP or SCP server based on your configuration. When you configure the appliance to save the capture file locally and later enable FTP or SCP, the appliance copies all the data starting with the oldest data. Infoblox recommends that you constantly monitor the FTP or SCP server to ensure that it has sufficient disk space. DNS queries and responses are stored on the appliance if the FTP or SCP server becomes unreachable. The maximum storage capacity varies based on the appliance model. After reaching the maximum limit, the appliance overwrites the old data with the new one. For information about the maximum hard drive space, see the table below. The amount of data captured depends on the DNS query rate and the domains that are included in or excluded from the capture. For information about how to exclude domains, see [Excluding Domains From Query and Response Capture](#) below.

You can also use the `dnstap` log format to achieve performance query logging. For information about `dnstap` implementation and configuring `dnstap`, see [Configuring dnstap](#).

### Capturing DNS Queries

You can capture queries to all domains or limit the capture to specific domains. You can also apply the Bulk Add Domains feature to tailor query capture to a desired subset of domains or zones. When capturing DNS queries, NIOS matches the specified domain name(s) and everything that belongs to the domain. For example, when you specify 'foo.com' as the domain, NIOS captures queries sent to 'foo.com,' 'mail.foo.com,' and 'ftp.foo.com.' NIOS captures queries to domains for which a name server is authoritative; it also captures recursive queries. Note that this feature does not support wildcard characters or regular expressions.

### DNS Query Message Format

The DNS query generates a query message in the following format:

```
<dd-mm-YYYY HH:MM:SS.uuu> <client IP>#<port> query: <query_Domain name>  
@0x7fbad80bda00 <class name> <type name> <- or +>[SETDC] <(name server ip)>
```

where

+ = recursion

- = no recursion

S = TSIG

E = EDNS option set

T = TCP query

D = EDNS 'DO' flag set

C = 'CD' message flag set

Following is a sample DNS query message:

```
30-Apr-2013 13:35:02.187 client 10.120.20.32#42386: query: foo.com  
@0x7fbad80bda00 IN A + (100.90.80.102)
```

## Capturing DNS Responses

You can capture DNS responses for the DNS queries sent to the server. The amount of data captured depends on the domains that are included in or excluded from the capture. A DNS response is based on a query generated for a domain. In the response message, NIOS captures the TTL value of a resource record, the resource record type, and resource data.

Following are characteristics of the response messages:

- They log only the answer section and do not include the authority and additional sections.
- Responses to all queries are logged, including queries with the type "ANY."
- The RR (resource record) list is not available at the end of a response message if rcode has a value other than NOERROR or if the response is NOERROR (nodata).
- Responses to all RR types, including those records not managed by NIOS such as HINFO records, are logged. However, there are few exceptions for some of the scenarios with DNSSEC records.
- Responses containing DNSSEC RRs (DNSKEY, DS, NSEC, NSEC3, NSEC3PARAM, RRSIG) when queried for non-DNSSEC RRs are not logged. However, responses are logged if a DNSSEC RR is explicitly queried.
- DNS updates are not logged in responses.

## DNS Response Message Format and Examples

The DNS query generates a response message in the following format:

```
<dd-mm-YYYY HH:MM:SS.uuu> <client IP>#<port> query: <query_Domain name> <class name>  
<type name> <- or +>[SETDC] <(name server ip)>
```

Flags = <- or +>[ATEDVL]

where

- = recursion not available

+ = recursion available (from DNS message header)

A = authoritative answer (from DNS message header)

t = truncated response (from DNS message header)

E = EDNS OPT record present (from DNS message header)

D = DNSSEC OK (from EDNS OPT RR)

V = responding server has validated DNSSEC records

L = response contains DTC synthetic record

The following is a sample DNS query message:

```
30-Apr-2020 13:35:02.187 client 10.120.20.32#42386: query: foo.com IN A +  
(100.90.80.102)
```

Following are some DNS response samples:

**Example 1: When querying an A record**

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: a2.foo.com IN A  
response: NOERROR +AED a2.foo.com. 28800 IN A 1.1.1.2;
```

**Example 2: When querying an AAAA record**

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: a4.foo.com IN  
AAAA response: NOERROR +AED a4.foo.com. 28800 IN AAAA ab::a;
```

**Example 3: When querying an A record over IPv6**

```
07-Apr-2013 20:16:49.083 client 2001::2#57398 UDP: query: a2.foo.com IN A  
response: NOERROR +AED a2.foo.com. 28800 IN A 1.1.1.2;
```

**Example 4: When querying an A record over TCP**

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 TCP: query: a2.foo.com IN A  
response: NOERROR +ED a2.foo.com. 28800 IN A 1.1.1.2;
```

**Example 5: When querying ANY record**

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: a2.foo.com IN  
ANY response: NOERROR +ED a2.foo.com. 28800 IN A 1.1.1.2;
```

**Example 6: When querying an A record with multiple addresses**

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: a1.foo.com IN A  
response: NOERROR +ED a1.foo.com. 28800 IN A 1.1.1.1; a1.foo.com. 28800 IN A  
11.1.1.1;
```

**Example 7: When querying an aliased A record**

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: c2.foo.com IN A  
response: NOERROR +ED c2.foo.com. 28800 IN CNAME a2.foo.com.; a2.foo.com. 28800  
IN A 1.1.1.2;
```

#### Example 8: When querying an NXDOMAIN

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: non-exist.foo.com IN A response: NXDOMAIN +ED
```

#### Example 9: Response message for NOERROR/nodata

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: a1.foo.com IN SRV response: NOERROR +ED
```

#### Example 10: Response message for refused query

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: refused.com IN A response: REFUSED +ED
```

#### Example 11: Response message when server fails

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#12345 UDP: query: servfail.com IN A response: SERVFAIL +E
```

#### Example 12: Response message when query A record in a signed zone

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: a1.signed.com IN A response: NOERROR +ED a1.signed.com. 28800 IN A 1.1.1.1;
```

#### Example 13: Response message for explicit query to DNSSEC RRs

```
07-Apr-2013 20:16:49.083 client 10.120.20.198#57398 UDP: query: a1.signed.com IN RRSIG response: NOERROR +ED a1.signed.com. 28800 IN RRSIG A 5 3 28800 20130616004903 20130611234903 4521 signed.com. evR0Ke7RbnkjFTsumT3JJg76bduFLfdEEEnszitXHQCbVYBS5rDy+qbUI HCQuN/ldCNTJbZQ8MEhuatzfms+2Y5K2sU67P9Yg6Gk0MxsT2LcJiBm/YqrYiZBWGKpLF6J0PdX05133Xwq8XxUStUEJxKfuzcKSY6jaSduQIdFL v6A=; a1.signed.com.900 IN RRSIG NSEC 5 3 900 20130616004903 20130611234903 4521 signed.com. CnFmXMx9D+ZkDsztQbW2xx8XCROGNMBp0baxFXS/Pxxhg4PQcq58laI97y2Xgqswn/wKNhY8p9hkes5+6t/ihCOIbw FryxtdivPfyYFf3jafedFN ymZu05K9bYUfCUZTGirzoJYhxBM7xFT8fMvxni9ngsbLym82Tqv3Nua 6wU=;
```

### Configuring DNS Query and Response Captures

To configure DNS query and response captures:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the **Toolbar** and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click **Toggle Advanced Mode** and select the **Logging** tab.
3. Under **Data Collection for all DNS Queries/Responses to a Domain**, complete the following:
  - Select the **Capture DNS Queries** checkbox to start capturing DNS queries. This enables the feature set for configuration. When you enable this option at the member level, the appliance captures DNS queries for the selected members only.

- Select the **Capture DNS Responses** checkbox to start capturing DNS responses. This enables the feature set for configuration. When you enable this option at the member level, the appliance captures DNS responses for the selected members only.



**Note**

Enabling the logging of queries and responses at the same time can increase disk space usage and adversely affect DNS services and performance. Infoblox recommends that you do not configure both logging of queries and logging of responses at the same time.

- Select **Capture queries/responses for all domains** to capture queries and responses to all domains and zones.
- Select **Limit capture to these domains** to capture DNS queries and responses to domains and zones one at a time.
- Specify domains for DNS capture operations in the Domain table by clicking the Add icon, and choosing **Add Domain** or **Bulk Add Domains** from the menu.
- To define the destination for capture files, do the following:
  - **Retain captured queries on the local disk:** Select this checkbox to save the DNS queries on the appliance. In addition to the local disk, you can select to export the DNS queries to the remote server by selecting **SCP** in the **Export to** drop-down list.
  - **Export to:** From the drop-down list, select **SCP** to back up the DNS queries on the remote server and **None** to save queries only on the appliance. To save the captured DNS queries on both the appliance and the remote server, select the **Retain captured queries on the local disk** checkbox and **SCP** from the **Export to** drop-down list. When you configure an SCP server and enable the MGMT port, the NIOS appliance uses SSH for data transfer. It uses the same authentication and provides the same security as SSH. SCP uses the LAN1 port to communicate with the external servers.
- When you select **FTP** or **SCP** from the **Export to** drop-down list, complete the following:
  - In the **Directory Path** field, enter the directory to which the capture file will be saved on the server. Infoblox recommends that you use the ~ symbol for the remote server.
  - In the **Server Address** field, enter the IP address of the remote server to which the capture files will be saved.
  - Enter the file server account **Username** and **Password** values.
- **Limit query data collected per file to minutes or 100MB (whichever comes first):** This option limits the collection of query data per capture file. A capture file for logging DNS queries and responses is rolled over based on the configured time limit or when the file reaches 100 MB in size, whichever is sooner. The default time limit is 10 minutes. You can enter a value from 1 to 10.

4. Save the configuration.

The following table lists the maximum hard drive space required for capturing DNS queries and responses for supported Infoblox appliance models.

*Maximum Hard Drive Space used for DNS queries and Responses*

Supported NIOS Appliances	Maximum Hard Drive Space for DNS Query/Response Capture (MB)
Trinzic 815 and IB-V815	900
Trinzic 825 and IB-V825	3100
Trinzic 1415 and IB-V1415	6000
Trinzic 1425 and IB-V1425	10000
Trinzic 2215 and IB-V2215	12000



Supported NIOS Appliances	Maximum Hard Drive Space for DNS Query/Response Capture (MB)
Trinzic 2225 and IB-V2225	28000
PT-1405	10000
PT-2205	28000

## Excluding Domains From Query and Response Capture

You can exclude individual domains and their subdomains from DNS query and response capturing. You can also use the Bulk Add Domains feature for a subset of domains to exclude them from query and response capturing. Subdomains can also be specified for exclusion. NIOS matches the specified domain names and their subdomains while filtering them in the Exclusion list. For example, when you specify 'foo.com' as the domain to be excluded, NIOS filters queries for 'foo.com,' 'mail.foo.com,' and 'ftp.foo.com.'



### Note

IDNs are not supported for the domains that are added to the Inclusion list and Exclusion list. You can use the punycode representation of an IDN in these lists.

To exclude a domain from query and response capturing, do the following:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the **Toolbar** and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab and click the **Members** tab -> *member* checkbox -> Edit icon.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click **Toggle Advanced Mode** and select the **Logging** tab.
3. Under **Data Collection for all DNS Queries/Responses to a Domain**, select the **Exclude the following domains** checkbox.
4. Click the Add icon and select **Add Domain** or **Bulk Add Domains** and specify domains in the Domain table.



### Note

NIOS first matches the domains in the Exclusion list and then matches the domains in the Inclusion list. NIOS does not capture queries and responses for the subdomains in the **Capture DNS Queries/Responses** list (Inclusion list) if their domains are added to the **Exclude the following domains list** (Exclusion list).

The following table provides examples of domains and subdomains added to the inclusion and exclusion lists and the corresponding effects on the query and response capture operations:

Capture DNS Queries/Responses (Inclusion List)	Exclude the Following Domains (Exclusion List)	Queried Domain	Captured Queries/Responses	Results
foo.com	it.foo.com	<ul style="list-style-type: none"> <li>foo.com</li> <li>finance.foo.com</li> </ul>	Yes	Does not match the exclusion list and therefore NIOS captures queries/responses made to foo.com and finance.foo.com.

Capture DNS Queries/ Responses (Inclusion List)	Exclude the Following Domains (Exclusion List)	Queried Domain	Captured Queries/ Responses	Results
		<ul style="list-style-type: none"> <li>it.foo.com</li> <li>ms.it.foo.com</li> </ul>	No	Matches the exclusion list and excludes their subdomains. NIOS does not capture queries/responses made to it.foo.com and ms.it.foo.com.
it.foo.com	foo.com			Domain is added to the exclusion list and its subdomain is added to the inclusion list. Therefore, this is not a valid configuration as queries/responses are not captured. The appliance displays a warning message for such invalid configuration.
it.foo.com	it.foo.com			Domain is added to both the exclusion and the inclusion lists. This is not a valid configuration as queries/responses are not captured. The appliance displays a warning message for such invalid configuration.
foo.com	corp1.com			Domain added to the inclusion list is not the subdomain of the domain added to the exclusion list. This is a redundant configuration as the outcome is the same even if the domain is removed from the Exclusion list. The appliance displays a warning message for such invalid configuration.
foo.com		<ul style="list-style-type: none"> <li>foo.com</li> <li>finance.foo.com</li> </ul>	Yes	Exclusion list is empty and therefore matches the Inclusion list. NIOS captures queries/responses made to foo.com and finance.foo.com
		<ul style="list-style-type: none"> <li>corp1.com</li> </ul>	No	NIOS does not capture queries/responses made to corp1.com as this domain is not mentioned in the inclusion list.
Capture All	foo.com	<ul style="list-style-type: none"> <li>foo.com</li> </ul>	No	Matches the exclusion list and NIOS does not capture queries made to foo.com.
		<ul style="list-style-type: none"> <li>finance.foo.com</li> </ul>	No	Subdomain matches the exclusion list and NIOS does not capture queries/responses made to finance.foo.com.
		<ul style="list-style-type: none"> <li>corp1.com</li> </ul>	Yes	Does not match the exclusion list. Matches the inclusion list and therefore NIOS captures queries/responses made to corp1.com.

## Monitoring Tools

You can use the audit log, the replication status, the traffic capture tool, and the capacity report in a Grid or HA pair to monitor administrative activities and capture traffic for diagnostic purposes. You can also use CLI commands to monitor certain DNS transactions.

This section includes the following topics:

- [Using the Audit Log](#)
- [Enabling Audit Log Rolling](#)
- [Specifying the Audit Log Type](#)
- [Viewing the Audit Log](#)
- [Searching in the Audit Log](#)
- [Downloading the Audit Log](#)
- [Viewing the Replication Status](#)

- [Using the Traffic Capture Tool](#)
- [Using the Capacity Report](#)
- [Monitoring DNS Transactions](#)
- [Viewing DNS Alert Indicator Status](#)
- [Configuring DNS Alert Thresholds](#)

In addition, if Grid members manage Microsoft servers, Grid Manager creates a synchronization log file for each managed Microsoft server. For information about viewing synchronization logs, see [Monitoring Managed Microsoft Servers](#).

## Using the Audit Log

The audit log contains a record of all Infoblox administrative activities. It provides the following detailed information:

- Timestamp of the change. If you have different admin accounts with different time zone settings, the appliance uses the time zone of the admin account that you use to log in to the appliance to display the date and timestamp.
- Administrator name.
- Changed object name.
- New value of the object. If you change multiple properties of an object, the audit log lists all changes in a comma-separated log entry. You can also search the audit log to find the new value of an object.

The appliance logs the following successful operations:

- Logins to Grid Manager and the API.
- Logout events that include: When users log out by clicking the **Logout** button, when the Grid Manager GUI times out, and when users are logged out due to an error.
- Write operations such as the addition, modification, and deletion of objects.
- System management operations such as service restarts and appliance reboots.
- Scheduled tasks such as adding an A record or modifying a fixed address.
- WAPI (RESTful API) session log information for each WAPI call, such as PUT, POST, and DELETE.

## Enabling Audit Log Rolling

When the audit log reaches its maximum size, which is 100 MB, the appliance automatically writes the file into a new file by adding a .0 extension to the first file and incrementing subsequent file extensions by 1. Files are compressed during the rotation process, adding a `.gz` extension following the numerical increment (`file.#.gz`). The sequential incrementation goes from zero through nine. When the eleventh file is started, the tenth log file (`file.9.gz`) is deleted, and subsequent files are renumbered accordingly. For example, the current log file moves to `file.0.gz`, the previous `file.0.gz` moves to `file.1.gz`, and so on through `file.9.gz`. A maximum of 10 log files (0-9) are kept. To list the audit log files and their sizes, log in to the Infoblox CLI and run the `show logfile` command.

To enable audit log rolling:

1. On the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Grid Properties** -> **Edit** in the Toolbar.
2. In the *Grid Properties* editor, on the **Security** tab, select **Enable Audit Log Rolling**.

## Specifying the Audit Log Type

Select **Detailed** (default), or **Brief**, or **WAPI Detailed** audit log type as follows:

1. On the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Grid Properties** -> **Edit** in the Toolbar.  
On the **System** tab, select the **System Manager** tab, and then click **System Properties** -> **Edit** in the Toolbar.
2. In the *Grid Properties Editor*, on the **General** tab, or in the *System Properties Editor*, on the **System** tab, select one of the following:

- **Detailed:** This is the default type. When you select this, Grid Manager displays detailed information on all administrative changes such as the timestamp of the change, administrator name, changed object name, and the new values of all properties in the logged message.
- **Brief:** Provides information on administrative changes such as the changed object name and action in the log message. The logged message does not show timestamp or admin name.
- **WAPI Detailed:** Select this option to view detailed RESTful API session information logs for successful WAPI calls such as PUT, POST, and DELETE. You can view the following session log information for each successful WAPI call:
  - **URI:** URI contains certain part of the incoming WAPI request. Example: version of WAPI and the associated object.
  - **InData:** InData contains input data fields of the incoming WAPI request. Example: Data field of the incoming WAPI request.
  - **Response Time:** Response time is calculated as the time difference between a WAPI request received and the response sent.

This option helps you to troubleshoot and monitor performance issues that impact specific WAPI calls and track WAPI usage. When you select this option, you can view additional columns such as **URI**, **InData** and **Response Time** in the *Audit log*.

The following example shows an audit log entry for a POST WAPI call:

```
[2018-05-29 09:20:12.026Z] [admin]: Created(POST) v2.9/zone_auth
{"fqdn":"foo.com"} 2.233 AuthZone foo.com DnsView=default: Set
fqdn="foo.com"
```

In the example above:

- `POST` indicates the WAPI call
- `v2.9/zone_auth` is the URI
- `{"fqdn":"foo.com"}` represents InData
- `2.233` is the response time.



#### Note

There might be a slight impact on the device performance as the session log information, such as URI, InData, and response time, are captured for all the successful WAPI calls. The audit log file size increases as the log entries increase, which may impact the storage capacity. Infoblox recommends that you select the **Copy Audit Log Messages to Syslog** checkbox in the *Grid Properties Editor* to move audit log information to the syslog and to an external server for longer retention. For more information about specifying syslog servers, see [Using a Syslog Server](#). All Cloud WAPI, via Cloud Network Automation (CNA) or Cloud Platform (CP) appliances, related events are logged to syslog instead of the audit log. For more information, see [Cloud Network Automation](#).

## Viewing the Audit Log

To view an audit log:

1. On the **Administration** tab, select the **Logs** tab -> **Audit Log** tab.
2. Optionally, use the filters to narrow down the audit log messages you want to view. Click **Show Filters** to enable the filters. Configure the filter criteria, and then click **Apply**.

Based on your filter criteria (if any), Grid Manager displays the following in the *Audit Log* viewer:

- **Timestamp:** The date, time, and time zone the task was performed. The time zone is the time zone configured on the member.

- **Admin:** The admin user who performed the task.  
Note that the admin user displayed as \$admin group name\$ represents an internal user. You can create an **admin** filter with “**matches expression**” equals `^[^$]` to filter out internal users.
- **Action:** The action performed. This can be CALLED, CREATED, DELETED, LOGIN\_ALLOWED, LOGIN\_DENIED, MESSAGE, MODIFIED, POST, PUT, and DELETE.
- **Object Type:** The object type of the object involved in this task. This field is not displayed by default. You can select this for display.
- **Object Name:** The name of the object involved in this task.
- **Execution Status:** The execution status of the task. Possible values are **Executed, Normal, Pending Approval** and **Scheduled**.
- **URI:** A certain part of the incoming WAPI request.
- **InData:** Input data fields of the incoming WAPI request.
- **Response Time:** The processing time of the incoming WAPI request.
- **Message:** Detailed information about the performed task.

You can also perform the following in the log viewer:

- Toggle between the single line view and the multi-line view for display.
- Navigate to the next or last page of the file using the paging buttons.
- Refresh the audit log view.
- Click the Follow icon to have the appliance automatically refresh the log every five seconds.
- Download the log.
- Clear the contents of the audit log.
- Use filters and the **Go To** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria.
- Export or print the content of the log.

## Searching in the Audit Log

Instead of paging through the audit log file to locate messages, you can have the appliance search for messages with certain text strings.

To search for specific messages, enter a search value in the search field below the filters, and then click the **Search** icon.

The appliance searches through the audit log file and highlights the search value in the viewer. You can use the arrow keys next to the Search icon to locate the previous or next message that contains the search value.

## Downloading the Audit Log

You can download the audit log file to a specified directory if you want to analyze it later. To download an audit log file:

1. On the **Administration** tab, select the **Logs** tab -> **Audit Log** tab, and then click the Download icon.
2. Navigate to a directory where you want to save the file, optionally change the file name (the default name is *auditLog.tar.gz*), and then click **OK**. If you want to download multiple audit log files to the same location, rename each downloaded file before downloading the next.



### Note

If your browser has a pop-up blocker enabled, you must turn off the pop-up blocker or configure your browser to allow pop-ups for downloading files.

## Viewing the Replication Status

The *Replication Status* panel reports the status of the database replication between Grid members and Grid Master, and between the two nodes in an independent HA pair. You can use this information to check the health of the Grid and HA pair activity.

To view the current replication status, from the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click

### Toggle Replication Status View.

Grid Manager can display the following replication information for each member:

- **Name:** The FQDN (fully qualified domain name) of the appliance.
- **Send Queue:** The size of the queue from the Grid Master to the Grid member.
- **Last Send:** The timestamp of the last replication information sent by the Grid Master.
- **Receive Queue:** The size of the queue from the Grid member to the Grid Master.
- **Last Receive:** The timestamp of the last replication information received by the Grid Master.
- **Member Replication Status:** The replication status between the member and the Grid Master. Grid Manager displays the status in green when the status is fine or red when the member is offline.
- **HA Replication Status:** The HA replication status between the active and passive nodes. The status is at the member level, not at the node level. Grid Manager displays the status in red when one of the nodes is offline.
- **Status:** The current operational status of the appliance. The status can be one of the following:
  - **Green:** The appliance is operating normally in a "Running" state.
  - **Yellow:** The appliance is connecting or synchronizing with its Grid Master.
  - **Red:** The Grid member is offline, is not licensed (that is, it does not have a DNSOne license with the Grid upgrade that permits Grid membership), is upgrading or downgrading, or is shutting down.
- **IPv4 Address:** The IPv4 address of the appliance or the VIP of an HA pair.
- **IPv6 Address:** The IPv6 address of the appliance or the VIP of an HA pair.
- **Identify:** This field appears only if your appliance has the unit identification button. This can be **On** or **Off**. When you identify the appliance by pressing the UID button on the appliance or through the GUI or the CLI command, this field displays **On**. Otherwise, this is **Off**.
- **DHCP, DNS, TFTP, HTTP, FTP, NTP, bloxTools, Captive Portal, DNS Accelerator Usage, Discovery, Reporting:** The current status of the service. The status can be one of the following:
  - **Green:** The service is enabled and running properly.
  - **Yellow:** The service is enabled, but there may be some issues that require attention.
  - **Red:** The service is enabled, but it is not running properly. A red status icon can also appear temporarily when a service is enabled and begins running, but the monitoring mechanism has not yet notified the Infoblox GUI.
  - **Gray:** The service is not configured or it is disabled.
- **Hardware Type:** The hardware type of the appliance, such as IB-1400.
- **Serial Number:** The serial number of the appliance.
- **DB Utilization:** The percentage of the database that is currently in use.
- **Comment:** Information about the appliance.
- **Site:** The location to which the member belongs. This is one of the predefined extensible attributes.
- **HA:** Indicates whether the member is an HA pair. If the member is an HA pair, Grid Manager displays the status of the HA pair.
- **Hardware Model:** The hardware model of the appliance.

You can do the following:

- Use filters and the **Go To** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria.
- Modify some of the data in the table. Double-click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).
- Edit the properties of a member.
  - Click the checkbox beside a member, and then click the Edit icon.
- Delete a member.
  - Click the checkbox beside a member, and then click the Delete icon.
- Export or print the list.

### Using the Traffic Capture Tool

You can capture the traffic on one or all of the ports on a NIOS appliance, and then view it using a third-party network protocol analyzer application, such as the Wireshark – Network Protocol Analyzer™. Using this tool, you can capture traffic for a single member or multiple Grid members simultaneously. The NIOS appliance must have a minimum of 500 MB of free disk space to start the traffic capture; otherwise, the traffic capture might fail.

The NIOS appliance saves all the traffic it captures in a .cap file and compresses it into a .tar.gz file. Your management

system must have a utility that can extract the .tar file from the .gz file, and an application that can read the .cap (capture) file format. The size of the .cap file is limited to 2 GB for Infoblox-4030-10GE, and the size is limited to 1 GB for all other NIOS appliances. You can also transfer the traffic capture file to your local management system, a TFTP server, an FTP server, or a SCP server.

NIOS saves failed traffic capture file transfer attempts in the syslog. Traffic capture file transfers can fail due to remote server issues, such as an invalid server IP address, or an incorrect username, or a password, or an invalid path. Note that the traffic capture fails when the disk space on the member is insufficient, but the file transfer will be successful. The NIOS appliance logs a warning message in the syslog when the traffic capture fails. You can find the following information about traffic capture and file transfers in the audit log:

- Start and stop actions performed on the members for traffic capture.
- Whether the traffic capture file was transferred to a server or downloaded to a local directory. For more information about the audit log, see [Using the Audit Log](#) above.

This section explains the process of capturing traffic, and how to download the traffic capture file to your management system. After that, you can extract the traffic capture file and view it with a third-party traffic analyzer application. The traffic capture file is shared between admin users.

You can also configure Grid Manager to trigger a traffic capture at set intervals and parameters. If Grid Manager detects that a parameter has breached a configured threshold or crossed the configured duration of time, it triggers a traffic capture. For more information about automated traffic capture, see [Enabling Automated Traffic Capture](#).



#### Note

The NIOS appliance always saves a traffic capture file as `<member name>_<timestamp>_tcpdumpLog.tar.gz`. Example: `infoblox.localdo_0_2018-10-15-03-47-53_tcpdumpLog.tar.gz`. For FTP and TFTP transfers, the name of the file is added as a prefix. Example: `filename.infoblox.localdo_0_2018-11-09-09-30-07_tcpdumpLog.tar.gz`

For a single member, you can also capture traffic on the NIOS appliance through the [Infoblox CLI](#) using the `set traffic_capture` command. However, you cannot use this command to capture traffic for multiple members. NIOS displays the traffic capture status and it allows you to download the captured traffic, irrespective of whether the traffic capture is initiated from the Infoblox CLI or from Grid Manager.

To capture traffic for a single member or multiple Grid members:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Traffic Capture** from the Toolbar.  
OR  
From the **Administration** tab, select the **Logs** tab → **Syslog** tab, and then click **Traffic Capture** from the Toolbar.
2. In the *Traffic Capture* dialog box, complete the following:  
**Members**
  - **Name:** Click the Add icon to add either a single or multiple Grid members for which you want to capture traffic. When you click the Add icon, Grid Manager displays the *Member Selector* dialog box from which you can select one or multiple members. Use SHIFT+click to select multiple contiguous rows or use CTRL+click to select multiple non-contiguous rows. Click **OK**. The selected members are added to the list of members in the **Members** table. You cannot add offline members to the list or capture traffic on an offline member.  
Selecting members in the **Grid Manager** → **Members** tab does not capture traffic for the selected member. To capture traffic, you must select members from the *Member Selector* dialog box.
  - **Interface:** Select the port on which you want to capture traffic. You can view the selected interface while the traffic capture is in progress. Note that if you enabled the LAN2 failover feature, the LAN and LAN2 ports generate the same output and Grid Manager displays the interface as **BOND** while the traffic capture is in progress. By default, the interface is set to **ALL** after the traffic capture process stops. For information about the LAN2 failover feature, see as described in [About Port Redundancy](#).

For vNIOS appliances, some of the options in the drop-down list may vary depending on your vNIOS configuration. For example, if you are using a single network interface instance of a vNIOS for GCP appliance, you will see choices specific to the LAN1 interface only.

- **LAN:** Select this to capture all the traffic the LAN port receives and transmits.



- **MGMT:** Select this to capture all the traffic the MGMT port receives and transmits.
  - **LAN2:** Select to capture all the traffic the LAN2 port (if enabled) receives and transmits.
  - **ALL:** Select this to capture the traffic addressed to all ports. Note that the NIOS appliance only captures traffic that is addressed to it.
  - **LANxnnnn:** If you have configured VLANs on the LAN1 or LAN2 port, the appliance displays the VLANs in the format LANx nnnn, where x represents the port number and nnnn represents the associated VLAN ID.
- **File Size:** Displays the size of the traffic capture log file, in kilobytes, for the respective member.
  - **Status:** Displays the status of the traffic capture session on the member. The status can be one of the following:
    - **STOPPED:** Indicates that the traffic capture session has stopped.
    - **RUNNING:** Indicates that the traffic capture session is in progress.
    - **NOT STARTED:** Indicates that the traffic capture session has not started.
  - **Transfer Status:** Displays the status of the traffic capture file transfer. The status can be one of the following:
    - **NOT STARTED:** Indicates that the file transfer has not started.
    - **STARTED:** Indicates that the file transfer has started.
    - **COMPLETED:** Indicates that the file transfer has been completed.
    - **FAILED:** Indicates that the file transfer has failed.
3. **Seconds to run:** Specify the number of seconds you want the traffic capture tool to run.
  4. **Capture Control:** Click the Start icon to start the capture. Note that the start action will overwrite the existing traffic capture file. You can click the Stop icon to stop the capture after you start it.
  5. **Transfer To:** Select the destination to transfer the traffic capture file. You can select **My Computer**, **TFTP**, **FTP**, or **SCP** from the drop-down list.
    - **My Computer:** Transfer the traffic capture file to a local directory on your computer. This is the default. Note to avoid consumption of the Grid Master disk space, NIOS restricts downloading the traffic capture files from multiple members to a local directory on your computer.
    - **TFTP:** Transfer the traffic capture file to a TFTP server.
      - **Filename:** Enter the directory path and the file name of the traffic capture file. For example, you can enter `/home/test/traffic_capture_filename` where `traffic_capture_filename` is the name of the file.
      - **IP Address of TFTP Server:** Enter the IP address of the TFTP server to which you want to transfer the traffic capture file.
    - **FTP:** Transfer the traffic capture file to an FTP server.
      - **Filename:** Enter the directory path and the file name of the traffic capture file. For example, you can enter `/home/test/traffic_capture_filename` where `traffic_capture_filename` is the name of the file.
      - **IP Address of FTP Server:** The IP address of the FTP server.
      - **Username:** Enter the username of your FTP account.
      - **Password:** Enter the password of your FTP account.
    - **SCP:** Transfer the traffic capture file to an SCP server.
      - **Filepath:** Enter the directory path of the traffic capture file. For example, you can enter `/home/test/`.
      - **IP Address of SCP Server:** The IP address of the SCP server.
      - **Username:** Enter the username of your SCP account.
      - **Password:** Enter the password of your SCP account.
  6. **Uncompressed Capture File Size:** Select the members for which you want to download the traffic capture file and then click **Download** to download the captured traffic. You can download and save the file only after the capture stops, but not when the tool is running. You can rename the file if you want. NIOS updates the size of the report when the capture tool is running.  
Note the NIOS appliance must have free disk space of at least 500MB + size of the traffic capture file (2 GB/1 GB, depending on the appliance model) to download the traffic capture file.
  7. **Last updated:** The timestamp of the last traffic capture process.
  8. Use terminal window commands (Linux) or a software application (such as StuffIt™ or WinZip™) to extract the contents of the .tar.gz file.

9. When you see the traffic.cap file in the directory where you extract the .tar.gz file, open it with a third-party network protocol analyzer application.

### Limitations of the Traffic Capture Tool

- You cannot add members to the list of members in the **Members** table on the *Traffic Capture* wizard when traffic capture is in progress.
- While the traffic capture file transfer is in progress on any member on the *Traffic Capture* wizard, Grid Manager does not allow you to start or stop traffic capture on members. However, you can start traffic capture using the CLI command even though the file transfer is in progress.
- If the traffic capture has already started on Grid Manager, you cannot initiate the capture again using the CLI command. You must wait until the process is complete.
- Traffic captures that are initiated using the CLI/WAPI commands do not appear in the **Members** table. To initiate traffic capture for a single member that was initiated and stopped using the CLI/WAPI command, add the respective member manually to the list of members in the **Members** table.
- Grid Manager does not allow you to initiate traffic capture on a list of selected members, if it has already been initiated for a member in the list using CLI/WAPI commands.

### Using the Capacity Report

You can view the capacity usage and object type information of an appliance in a capacity report. The capacity report displays capacity and object type information of an independent appliance, a Grid Master, or a Grid member. For an HA pair, the report displays information on the active node.

The top half of the panel displays a capacity summary, and the bottom half displays the object types the appliance supports and the total counts for each object type.

To view a capacity report:

- From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox, and then click **Capacity Report** from the Toolbar.

The capacity summary contains the following information:

- **Name:** The name of the appliance.
- **Role:** The role of the appliance. The value can be **Grid Master**, **Grid Master Candidate**, **Grid Member**, or **Standalone**.
- **Hardware Type:** The type of hardware. For an HA pair, the report displays the hardware type for both the active and passive nodes.
- **Object Capacity:** The maximum number of objects the appliance can support.
- **Total Objects:** The total number of objects currently in the database.
- **% Capacity Used:** The percentage of the capacity in use.

The capacity report filters object types you can manage through the appliance. You can configure the object types you want to see in the following table by completing the following in the **Minimum Object Total** filter:

- **Minimum Object Total:** Enter the minimum number of objects within an object type of which Grid Manager displays. In the Object Type table, Grid Manager displays only the object types that contain at least the specified number of objects you enter in this field.

The capacity report displays the following information:

- **Object Type:** The type of objects. For example, DHCP Lease, Admin Group, or PTR Record. For objects that are only used for internal system operations, the report groups and shows them under **Other**.
- **Total:** The total number of objects for the specific object type. You can print the object type information or export it to a CSV file.

### Monitoring DNS Transactions

The NIOS appliance provides tools for monitoring DNS transactions and mitigating cache poisoning from UDP (User Datagram Protocol) traffic on source port 53. Cache poisoning can occur when a DNS server accepts maliciously created unauthentic data. The DNS server ends up locally caching the incorrect entries and serving them to users that make the same DNS requests. In a maliciously created situation, the attacker can redirect Internet traffic from the legitimate host to

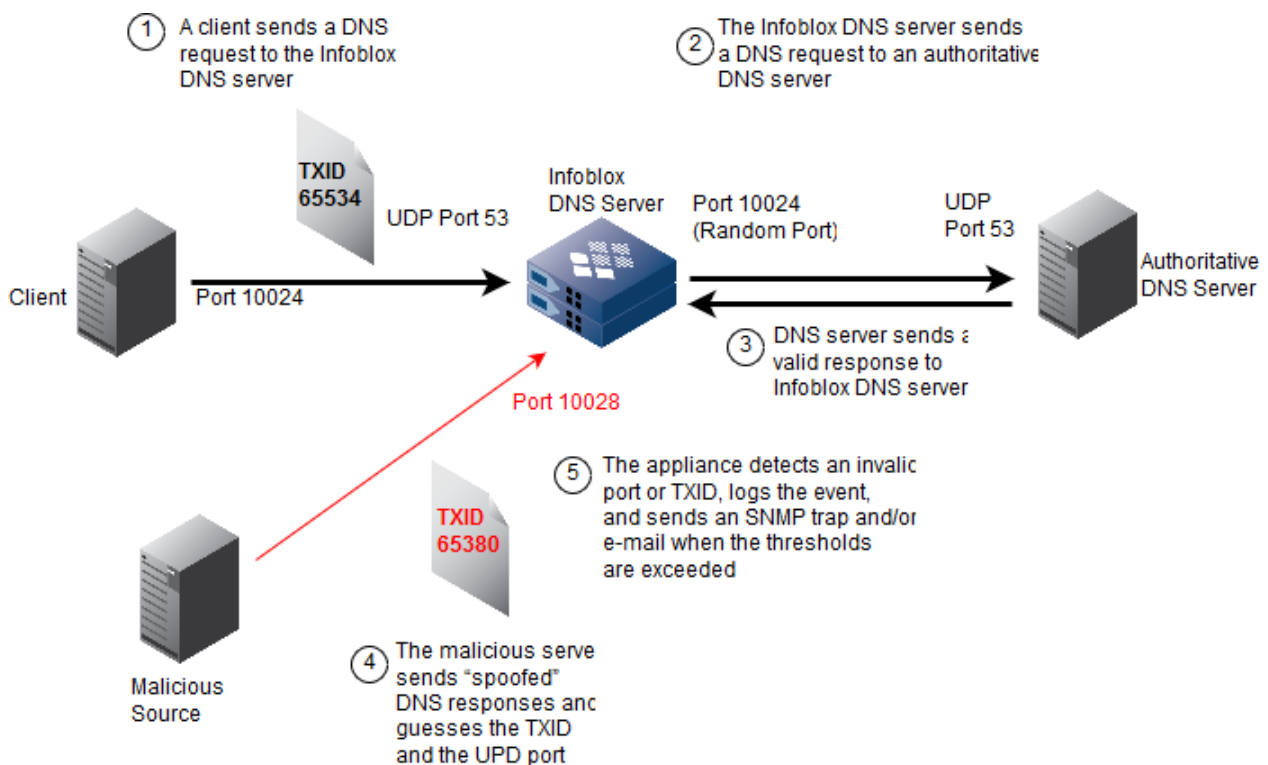
another host that the attacker controls.

You can configure the appliance to track invalid DNS responses for recursive DNS queries. The appliance tracks DNS responses that arrive on invalid ports or have invalid TXIDs (DNS transaction IDs). Both invalid ports and invalid TXIDs could be indicators of cache poisoning. An invalid port is a DNS response that arrives from UDP (User Datagram Protocol) port 53 with either one of the following conditions:

- There are no outstanding DNS requests from the port on which the response arrives.
- The TXID of the DNS response matches the TXID of an outstanding request. However, the request was sent from a port other than the port on which the response arrives.

An invalid TXID is a DNS response that arrives from UDP port 53, and the TXID does not match the TXID of an outstanding DNS request. The below figure illustrates how the appliance detects an invalid port and an invalid TXID.

#### Invalid Port and Invalid TXID



Both invalid ports and invalid TXIDs could be indicators of DNS cache poisoning, although a small number of them is considered normal in situations where valid DNS responses arrive after the DNS queries have timed out. You can configure the appliance to track these indicators, and you can view their status. You can also configure thresholds for them. When the number of invalid ports or invalid TXIDs exceeds the thresholds, the appliance logs an event in the syslog file and sends an SNMP trap and e-mail notification, if you enable them. You can then configure rate limiting rules to limit incoming traffic or completely block connections from primary sources that send the invalid DNS responses.

Rate limiting is a token bucket system that accepts packets from a source based on the rate limit. You can configure the number of packets per minute that the Infoblox DNS server accepts from a specified source. You can also configure the number of packets for burst traffic, which is the maximum number of packets that the token bucket can accept. Once the bucket reaches the limit for burst traffic, it discards the packets and starts receiving new packets according to the rate limit.

The appliance monitors only UDP traffic from remote port 53 for the following reasons:

- The attacks that the appliance monitors do not happen over TCP.
- DNS responses are sent only from port 53. The appliance discards DNS responses that are sent from other ports.

To monitor invalid ports and invalid TXIDs on the Infoblox DNS server, follow these procedures:

- Enable DNS network monitoring and DNS alert monitoring. For information, see [Enabling and Disabling DNS Alert Monitoring](#) below.
- Configure the thresholds for DNS alert indicators. For information, see [Configuring DNS Alert Thresholds](#) below.
- Enable SNMP traps and e-mail notifications. For information, see [Configuring SNMP](#).
- Review the DNS alert status. For information, see [Viewing DNS Alert Indicator Status](#) below.
- Identify the source of the attack by reviewing the DNS alert status, syslog file, and SNMP traps. For information on SNMP traps for DNS alerts and threshold crossing traps, see [SNMP MIB Hierarchy](#).

To mitigate cache poisoning, you can limit incoming traffic or completely block connections from specific sources, as follows:

- Enable rate limiting on the DNS server. For information, see [Enabling and Disabling Rate Limiting from External Sources](#) below.
- Configure rate limit traffic rules from specific sources. For information, see [Configuring Rate Limiting Rules](#) below.

You can verify the rate limiting rules after you configure them. For information, see [Viewing Rate Limiting Rules](#) below.

### Enabling and Disabling DNS Alert Monitoring

The appliance monitors only UDP traffic on port 53 for recursive queries, and then reports invalid DNS responses. DNS alert monitoring is disabled by default. For an HA pair, you must enable DNS alert monitoring on both the active and passive nodes.

To enable DNS network monitoring and DNS alert monitoring:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
set monitor dns on
```

The appliance displays the following:

```
Turning on DNS Network Monitoring...
```

3. Enter the following command:

```
set monitor dns alert on
```

When you enable DNS alert monitoring, if DNS network monitoring is disabled, the appliance automatically enables DNS network monitoring and displays the following:

```
DNS Network Monitoring is disabled. It must be enabled for alerting to function.
```

```
Enable DNS Monitoring now? (y or n):
```

You can also disable DNS network monitoring and DNS alert monitoring using the following commands:

```
set monitor dns off
```

```
set monitor dns alert off
```



#### Note

When you restart DNS network monitoring, you also reset the SNMP counters for DNS alerts.

You can then view the alert status to identify the primary source of invalid DNS responses.

### Viewing DNS Alert Indicator Status

To view DNS alert indicator status:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
show monitor dns alert status
```

The appliance displays historical alert counts and up to five primary sources that generate invalid DNS responses, as shown in the following example:

```
Data last updated: Mon Oct 6 14:47:12 2008
```

```
DNS Alert1m5m15m60m24hEver
```

```
=====
```

```
port8 12 1212 12 12
```

```
txid8 12 1212 12 12
```

```
There were 80 DNS responses seen in the last minute.
```

```
10% were to an invalid port.
```

```
10% had an invalid TXID.
```

```
Primary sources of invalid responses:
```

```
4.4.4.4 (unknown) sent 4
```

```
2.2.2.2 (unknown) sent 3
```

```
7.7.7.7 (unknown) sent 1
```

The appliance attempts to resolve the hostnames of the sources that sent invalid responses, if the DNS resolver is enabled. If the appliance cannot resolve a hostname, it displays "unknown" as the hostname of the invalid response.

### Configuring DNS Alert Thresholds

You can configure thresholds for DNS alerts to control when the appliance tracks DNS attacks on UDP port 53 and issues SNMP traps and e-mail notifications.



#### Note

Ensure that you enable SNMP traps and e-mail notifications. For information, see [Configuring SNMP](#).

You can configure thresholds for both invalid ports and invalid TXIDs. The default thresholds for both invalid ports and TXIDs are 50%. When the number of invalid ports or invalid TXIDs exceeds the thresholds, the appliance logs the event and sends SNMP traps and notifications. You can configure the thresholds either as absolute packet counts or as percentages of the total traffic during a one-minute time interval.

To configure DNS alert thresholds:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
set monitor dns alert modify port | txid over threshold_value packets |  
percent  
where
```

`port | txid` = Enter `port` to set the threshold for invalid ports, or enter `txid` to set the threshold for invalid TXIDs.

`threshold_value` = Enter the number of packets or percentage for the threshold.

`packets | percent` = Enter `packets` if you want to track the total packet count, or enter `percentage` if you want to track a percentage of the total traffic. For a percentage-based threshold, the appliance does not generate a threshold crossing event if the traffic level is less than 100 packets per minute.

For example, if you want the appliance to send a DNS alert when the percentage of DNS responses arriving on invalid ports from UDP port 53 exceeds 70% per minute, you can enter the following command:

```
set monitor dns alert modify port over 70 percent
```

If you want the appliance to send a DNS alert when the total number of packets with invalid TXIDs from UDP port 53 is over 100 packets per minute, you can enter the following command:

```
set monitor dns alert modify txid over 100 packets
```

When there is a DNS alert, the appliance logs an event in the syslog file and sends an SNMP trap and e-mail notification if enabled.

### Viewing DNS Alert Thresholds

You can view the DNS alert thresholds. The appliance displays the current thresholds. If you have not configured new thresholds, the appliance displays the default thresholds, which are 50% for both invalid port and TXID.

To view the DNS alert thresholds:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
show monitor dns alert
```

The appliance displays the threshold information as shown in the following example:

```
DNS Network Monitoring is enabled. Alerting is enabled.
```

```
DNS Alert Threshold (per minute)
```

```
=====
```

```
portover 70% of packets
```

```
txidover 100 packets
```

### Enabling and Disabling Rate Limiting from External Sources

You can mitigate cache poisoning on your DNS server by limiting the traffic or blocking connections from UDP port 53. To enable rate limiting from sources:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
set ip_rate_limit on
```

The appliance displays the following:

```
Enabling rate limiting will discard packets and may degrade performance.
```

```
Are you sure? (y or n):
```

When you enable rate limiting, the appliance discards packets based on the configured rate limiting rules. This might affect the DNS performance when the appliance discards valid DNS responses.

3. Enter **y** to enable rate limiting.

When you enable rate limiting, the appliance applies the rate limiting rules that you configured. You might want to configure the rate limiting rules before enabling rate limiting. For information on how to configure rate limiting rules, see [Configuring Rate Limiting Rules](#) below.

You can also disable rate limiting by entering the following command:

```
set ip_rate_limit off
```

When you disable rate limiting, the appliance stops applying the rate limiting rules.

### Configuring Rate Limiting Rules

You configure rate limiting rules to limit access or block connections from UDP port 53. The rules take effect when you enable rate limiting.

When adding rules, ensure that you do not include an IP address that matches the IP address of either the Grid Master or Grid member. Doing this could affect VPN connectivity. To configure rate limiting rules:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
set ip_rate_limit add source all | ip_address [/mask] limit packets/m [burst  
burst_packets]
```

where

`all | ip_address` = Enter `all` or `0.0.0.0` if you want to limit all traffic from all sources, or enter the IP address from which you want to limit the traffic.

`[/mask]` = Optionally, enter the netmask of the host from which you want to limit the traffic.

`packets` = Enter the number of packets per minute that you want to receive from the source.

`[burst burst_packets]` = Optionally, enter `burst` and the number of packets for burst traffic. This is the maximum number of packets accepted.

The following are sample commands and descriptions for rate limiting rules:

- To block all traffic from host 10.10.1.1, enter the following command:

```
set ip_rate_limit add source 10.10.1.1 limit 0
```

- To limit traffic to five packets per minute from host 10.10.1.2, enter the following command:

```
set ip_rate_limit add source 10.10.1.2 limit 5/m
```

- To limit the traffic to five packets per minute from host 10.10.2.1/24 with an allowance for burst traffic of 10 packets, enter the following command:

```
set ip_rate_limit add source 10.10.2.1/24 limit 5/m burst 10
```

- To limit the traffic to 5000 packets per minute from all sources, enter the following command:

```
set ip_rate_limit add source all limit 5000/m
```



## Removing Rate Limiting Rules

You can remove the existing rate limiting rules that limit access or block connections from UDP port 53. To remove all the existing rules:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:
  - To remove the rate limiting rule that limits traffic from all sources, enter:

```
set ip_rate_limit remove source all
```

or

- To remove all of the rate limiting rules from all sources, enter:

```
set ip_rate_limit remove all
```

To remove one of the existing rules for an existing host:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
set ip_rate_limit remove source ip-address[/mask]
```

## Viewing Rate Limiting Rules

You can view the existing rate limiting rules that limit access or block connections from UDP port 53. To view rate limiting rules:

1. Log in to the Infoblox CLI as a superuser account.
2. Enter the following CLI command:

```
show ip_rate_limit
```

The appliance displays the rules, as shown in the following example:

```
IP rate limiting is enabled.
```

Source	Limit	Burst
=====		
10.10.1.1	0 packets/minute	0 packets
10.10.1.2	5 packets/minute	5 packets
10.10.2.1/24	5 packets/minute	10 packets
all	5000packets/minute	5000 packets

\* Note: Copy audit to syslog feature should not alter. But increases the syslog file size.

## Using a Syslog Server

Syslog is a widely used mechanism for logging system events. NIOS appliances generate syslog messages that you can view through the Syslog viewer and download to a directory on your management station. In addition, you can configure a NIOS appliance to send the messages to one or more external syslog servers for later analysis. Syslog messages provide information about appliance operations and processes. NIOS appliances include syslog messages generated by the BloxTools service. You can choose logging categories to send specific syslog messages. The prefixes in the syslog messages are based on the logging categories you configure in the syslog. Note that syslog messages are prefixed only when you select logging categories. For information about how to configure logging categories, see [Specifying Syslog](#)

Servers below. You can also include audit log messages and specific BIND messages among the messages the appliance sends to the syslog server.

In addition to saving system messages to a remote syslog server, a NIOS appliance also stores the system messages locally. When the syslog file reaches its maximum size, which is 300 MB for Infoblox appliances and VMware virtual appliances, the appliance automatically writes the file into a new file by adding a `.0` extension to the first file and incrementing subsequent file extensions by 1.

Files are compressed during the rotation process, adding a `.gz` extension following the numerical increment ( `file.#.gz` ). The sequential incrementation goes from zero through nine. When the eleventh file is started, the tenth log file ( `file.9.gz` ) is deleted, and subsequent files are renumbered accordingly. For example, the current log file moves to `file.0.gz`, the previous `file.0.gz` moves to `file.1.gz`, and so on through `file.9.gz`. A maximum of 10 log files (0-9) are kept.

You can set syslog parameters for RPZ at the Grid, member, and zone levels. At the member level, you can override Grid-level syslog settings and enable syslog proxy, also you can override Grid-level settings to zone level. For more information see about Modifying RPZs, see [Managing RPZs](#).

You can configure the appliance to back up rotated syslog files to external servers through FTP or SCP. When you do so, the appliance forwards the rotated syslog files to the external servers that you configure. You can configure up to 10 external syslog backup servers each at the Grid, member, and zone levels. You can also override the Grid-level server configuration at the member level. For information about configuring syslog backup servers, see [Configuring Syslog Backup Servers](#) below.

This section includes the following topics:

- [Specifying Syslog Servers](#)
  - [Syslog Message Prefixes](#)
  - [IP Address Used in the Syslog Configuration File](#)
- [Configuring Syslog Backup Servers](#)
- [Configuring Syslog for Grid Members](#)
- [Setting DNS Logging Categories](#)
- [Viewing the Syslog](#)
- [Viewing the RPZ Threat Details](#)
- [Searching in the Syslog](#)
- [Downloading the Syslog File](#)

## Specifying Syslog Servers

To configure a NIOS appliance to send messages to a syslog server, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties* editor, select the **Monitoring** tab, and then complete the following:
  - **Syslog:** In addition to storing the syslog on a Grid member, you can configure the Grid to send the log to an external syslog server.
  - **Syslog size (MB):** Specify the maximum size for a syslog file. Enter a value between 10 and 300. The default is 300.  
When the syslog file reaches the size you enter here, the appliance automatically writes the file into a new file by adding a `.0` extension to the first file and incrementing subsequent file extensions by 1.
  - **Log to External Syslog Servers:** Select this to enable the appliance to send messages to a specified syslog server. Grid Manager displays the current syslog servers in the table. To define a new syslog server, click the Add icon and complete the following:
    - **Address:** Enter the IP address of the syslog server. Entries may be an IPv4 or IPv6 address.

- **Transport:** From the drop-down list, select whether the appliance uses **Secure TCP**, **TCP**, or **UDP** to connect to the external syslog server.
- **Server Certificate:** Click **Select** to upload a self-signed or a CA-signed server certificate. In the *Upload* dialog, click **Select** and navigate to the certificate file, and then click **Upload**. Note that this is valid only for **Secure TCP** transport.
- **Interface:** From the drop-down list, select the interface through which the appliance sends syslog messages to the syslog server.
  - **Any:** The appliance chooses any port that is available for sending syslog messages.
  - **LAN:** The appliance uses the LAN1 port to send syslog messages.
  - **MGMT:** The appliance uses the MGMT port if it has been configured. Otherwise, it uses the LAN1 port.
- **Source:** From the drop-down list, select which syslog messages the appliance sends to the external syslog server:
  - **Any:** The appliance sends both internal and external syslog messages.
  - **Internal:** The appliance sends syslog messages that it generates.
  - **External:** The appliance sends syslog messages that it receives from other devices, such as syslog servers and routers.
- **Node ID:** Specify the host or node identification string that identifies the appliance from which syslog messages are originated. This string appears in the header message of the syslog packet. Select one of the following:
  - **LAN:** Use the LAN1 IP address of the appliance. For an HA pair, this is the LAN1 address of the active or passive node. This is the default.
  - **Host Name:** Use the host name of the appliance in FQDN format.
  - **IP and Host Name:** Use both the FQDN and the IP address of the appliance. The IP address can be the LAN1 or MGMT IP address depending on whether the MGMT port has been configured. Note that if the MGMT port is not configured, the LAN1 IP address is used.
  - **MGMT:** Use the MGMT IP address, if the port has been configured. If the MGMT port is not configured, the LAN1 IP address is used. This can be an IPv4 or IPv6 address.
- **Port:** Enter the destination port number. The default is 514 for TCP and UDP. For Secure TCP, the default port is 6514.
- **Severity:** Choose a severity filter from the drop-down list. When you choose a severity level, the appliance sends log messages with the selected level and the levels above it. The severity levels range from the lowest, **debug**, to the highest, **emerg**. For example, if you choose **debug**, the appliance sends all syslog messages to the server. If you choose **err**, the appliance sends messages with severity levels **err**, **crit**, **alert**, and **emerg**.
  - **emerg:** Panic or emergency conditions. The system may be unusable.
  - **alert:** Alerts, such as NTP service failures, that require immediate actions.
  - **crit:** Critical conditions, such as hardware failures.
  - **err:** Error messages, such as client update failures and duplicate leases.
  - **warning:** Warning messages, such as missing keepalive options in a server configuration.
  - **notice:** Informational messages regarding routine system events, such as "starting BIND".
  - **info:** Informational messages, such as DHCPACK messages and discovery status.
  - **debug:** Messages that contain information for debugging purposes, such as changes in the latency timer settings and AD authentication failures for specific users.
- **Logging Category:** Select one of the following logging categories:
  - **Send all:** Select this to log all syslog messages, irrespective of categories to which it belongs. When you select this option, the appliance logs syslog messages for all the events, including all DNS and Infoblox related events. However, the syslog messages are not prefixed when you select this option.
  - **Send selected categories:** Select this to configure logging categories from the list of available logging categories. Use the arrows to move logging categories from the **Available** table to the **Selected** table and vice versa. The appliance sends syslog messages for the categories that are in the **Selected** table. When you select this option, you must add at least one logging category. The syslog messages are prefixed with a category name to which it belongs. Also, the RPZ events logged in the syslog messages uses specific prefixes for the selected categories. Note that the syslog messages are prefixed when you set logging categories for at least one external syslog server, even if you set other external

syslog servers as **Send All**. For more information about syslog prefixes, see Syslog Message Prefixes below.

- Click **Add** to add the external syslog server information.
  - Optionally, click the **Test** button to test the connection to the syslog server that you configured.
  - **Copy Audit Log Messages to Syslog**: Select this for the appliance to include audit log messages it sends to the syslog server. This function can be helpful for monitoring administrative activities on multiple appliances from a central location.
    - **Syslog Facility**: This is enabled when you select **Copy audit log messages to syslog**. Select the facility that determines the processes and daemons from which the log messages are generated.
3. Save the configuration and click **Restart** if it appears at the top of the screen.



#### Note

The syslog categories you specify here are different from the logging categories specified in the **Logging** tab in the *Grid DNS Properties* or *Member DNS Properties* editor. The external server preserves contents of the selected categories even when selection is changed from **Send all** to **Send selected categories** and vice versa.

## Syslog Message Prefixes

You can configure the syslog external backup servers to send (archive) syslog files to different destinations by their logging categories. This allows you to split syslog files based on the service and efficiently perform troubleshooting. For example, you can archive all DNS related logs on Server 1, and all DHCP related logs on Server 2. For information about how to configure an external syslog backup server, see *Specifying Syslog Servers* above.

When you select the **Send selected categories** option, the syslog messages are prefixed with a category name to which it belongs.

For syslog message prefixes to be enabled, you must check the **Log to External Syslog Servers** checkbox in **Grid Properties > Monitoring**. Also, the external syslog server (which can be a virtual or a physical server) must have at least one of the syslog categories selected instead of the **Send all** option selected in the **Logging Category** field.



#### Note

When you set **Send all** in the **Logging Category**, the appliance logs syslog messages for all the events and they are not prefixed. The syslog messages are prefixed even if one external syslog server is set with the **Send selected categories** option.

The following are the prefixes used for different logging categories:

- **DNS Logging Categories**: All DNS related messages use the following prefixes: `client`, `config`, `database`, `dnssec`, `general`, `lame_servers`, `network`, `notify`, `queries`, `query_rewrite`, `resolver`, `responses`, `rpz`, `security`, `update`, `update_security`, `xfer_in`, and `xfer_out`.

Sample syslog message for queries:

```
2022-01-18T09:35:35+00:00 daemon member1.com named[19355]: info client
@0x7fea340ccc90 10.111.45.104#34670 (a.solo.com): query: a.solo.com IN A
+E(0)K (10.34.122.22)
```

Sample syslog message for xfer-out:

```
2014-10-10T06:44:09+00:00 daemon infoblox.localdomain named[17630]: info
xfer-out:
```

```
client 10.120.20.157#58275 (zone.com): transfer of 'zone.com/IN': AXFR started
```

- **ADP:** All Infoblox related messages use prefix `adp`.

Note there is no prefix for RPZ syslog messages that do not belong to the DNS or ADP category.

- **DHCP:** All DHCP related messages use the following prefixes: `dhcpcd`, `omshell`, `dhcrelay`, and `dhclient`.

Sample syslog message for dhcp:

```
Sep 4 09:23:44 10.34.6.28 dhcpcd[20310]: DHCPACK on 70.1.20.250 to fc:5c:fc:5f:10:85 via
```

```
eth1 relay 10.120.20.66 lease-duration 600
```

- **DTC:** All DTC related messages use the following prefixes: `idns_healthd` and `idnsd`.

Sample syslog message for idns\_healthd:

```
Sep 3 12:12:35 10.34.6.30 idns_healthd[1220]: resource health status [Monitor 'icmp'
```

```
(ICMP, port 0) checked server 's1' (IP 10.34.6.23), status: IPv4=ONLINE]
```

- **Cloud:** All cloud related messages use prefix `cloud_api`.

Sample syslog message for cloud\_api:

```
Sep 4 10:53:30 10.34.6.32 cloud_api[5354]: [admin]: Login_Allowed - - to=Serial\040Console apparently_via=Remote ip=10.120.20.66 auth=Local group=.admin-group
```

- **NTP:** All NTP related messages use prefix `ntpd`.

Sample syslog message for NTP:

```
Sep 28 06:57:21 10.35.116.7 ntpd[12186]: precision = 0.053 usec
```

```
Sep 28 06:57:21 10.35.116.7 ntpd[12186]: Listening on interface #0 wildcard, 0.0.0.0#123
```

```
Disabled
```

- **File Distribution:** All File Distribution related messages use the following prefixes: `ftpd` and `tftp`.

Sample syslog message for TFTP:

```
Sep 3 13:03:09 10.34.6.30 monitor[23623]: Type: TFTP, State: Red, Event: A  
TFTPD daemon
```

```
failure has occurred
```

- **Authentication:** All Authentication related messages use the following prefixes: `auth`, `authpriv`, `AD`, and `radiusd`.

Sample syslog message for RADIUS authentication:

```
Sep 28 10:09:55 10.35.116.4 httpd: 2015-09-28 10:09:55.912Z [user1]:  
Login_Allowed - -
```

```
to=AdminConnector ip=10.120.253.227 auth=RADIUS group=admin-group  
apparently_via=GUI
```

- **Microsoft Integration:** All Microsoft Integration related messages use the following prefixes: `dns_server`, `connect_status`, `dns_zone`, `dhcp_server`, `dhcp_leases`, `clear_lease`, `ad_site`, and `ad_users`.

Sample syslog message for microsoft integration:

```
dns_server:
```

```
Sep 7 09:46:17 10.34.22.20 mssyncd[22315]: dns_server address  
10.102.30.157 : Conflict
```

```
in property Forwarders: NIOS value (property=<NULL IP array>) and Microsoft  
value
```

```
(property={10.0.2.35, 10.0.2.60}). Resolved by using the Microsoft value
```

```
dhcp_server:
```

```
Sep 7 10:08:48 10.34.22.20 mssyncd[22316]: dhcp_server address  
10.102.30.157 : Couldn't
```

```
open RPC interface <MS-WKST>: an instance of a named pipe cannot be found in  
the listening
```

```
state
```

```
Sep 7 10:08:48 10.34.22.20 mssyncd[22317]: dns_server address  
10.102.30.157 : Opened
```

```
RPC interface <MS-WKST> as user 'ad-15\frtest'
```

## IP Address Used in the Syslog Configuration File

The following table describes which IP address the appliance uses as the node ID in the syslog configuration file, provided that the MGMT port has been configured. If the MGMT port is not configured, the LAN1 IP address is always used regardless of the configuration.

### *IP address Used in Syslog Config File when MGMT Port is Configured*

Interface	Node ID	IP used in syslog configuration file
Any	MGMT	MGMT IP address
Any	IP and Host Name	MGMT IP address
MGMT	MGMT	MGMT IP address
MGMT	IP and Host Name	MGMT IP address
LAN	MGMT	LAN1 IP address
LAN	IP and Host Name	LAN1 IP address

## Configuring Syslog Backup Servers

You can configure external syslog backup servers to forward rotated syslog files. You can configure up to 10 external syslog backup servers.

To configure external backup servers, complete the following:

1. **Grid:** From the **Grid** tab -> **Grid Manager** tab, expand the Toolbar and click **Grid Properties** -> **Edit**.  
**Member:** From the **Grid** tab -> **Grid Manager** tab, click the **Members** tab, select the *member* checkbox, and click the **Edit** icon.
2. **Grid:** In the *Grid Properties* editor, select the **Syslog Backup** tab.  
**Member:** In the *Grid Member Properties* editor, select the **Syslog Backup** tab and then click **Override** to override the Grid-level settings.

To modify backup server settings, complete the following:

- **Address:** Enter the IP address of the external backup server. You are not allowed to configure more than one server using the same IP address at the same level (Grid or member). However, you can use the same server IP address at different levels (Grid or member). Note that you cannot modify the IP address for the overridden server.
  - **Protocol:** Select **SCP** or **FTP** from the drop-down list.
  - **Port:** Enter the destination port number. The default port is 20 for FTP and 22 for SCP.
  - **Path:** Enter the directory path for the syslog file.
  - **Username:** Enter the username of your FTP or SCP account.
  - **Password:** Enter the password of your FTP or SCP account. If you do not change the password of the overridden server, then make sure that you use the same password specified at the Grid level.
  - **Enabled:** Select this checkbox to enable the FTP or SCP server. The appliance forwards the rotated syslog files to the external servers that you configure only after you select this checkbox. Clear the checkbox to disable the server.
3. Click **Save and Close**.



## Configuring Syslog for Grid Members

You can override Grid-level syslog settings and enable syslog proxy for individual members. When you enable syslog proxy, the member receives syslog messages from specified devices, such as syslog servers and routers, and then forwards these messages to an external syslog server. You can also enable appliances to use TCP for sending syslog messages. Using TCP is more reliable than using UDP; this reliability is important for security, accounting, and auditing messages sent through the syslog. Note that you cannot enable syslog proxy for Grid members, if they are configured on a Grid Master.

To configure syslog parameters for a member, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.
2. In the *Grid Member Properties* editor, select the **Monitoring** tab -> **Basic** tab, click **Override** in the Syslog section, and then complete the fields as described in Specifying Syslog Servers above.  
In addition to storing the system log on a Grid member, you can configure a member to send the log to a syslog server.
3. Select the **Advanced** tab and complete the following:
  - **Enable syslog proxy:** Select this to enable the appliance to receive syslog messages from other devices, such as syslog servers and routers, and then forward these messages to an external syslog server.
    - **Enable listening on TCP:** Select this if the appliance uses TCP to receive messages from other devices. Enter the number of the port through which the appliance receives syslog messages from other devices.
    - **Enable listening on UDP:** Select this if the appliance uses UDP to receive messages from other devices. Enter the number of the port through which the appliance receives syslog messages from other devices.
  - **Proxy Access Control:** Select one of the following to configure access control when receiving syslog messages from specific syslog servers or routers:
    - **None:** Select this if you do not want to configure syslog proxy. When you select this option, none of the devices can send syslog messages to the appliance. This is selected by default.
    - **Named ACL:** Select this and click **Select Named ACL** to select a named ACL that contains only IPv4 addresses and networks. This does not support Tsig key based ACEs. When you select this, the appliance permits clients that have **Allow** permission in the named ACL to allow syslog messages from specific syslog servers or routers. You can click **Clear** to remove the selected named ACL.
    - **Set of ACLs:** Select this to configure individual access control entries (ACEs). Click the Add icon and select one of the following from the drop-down list. Grid Manager adds a row to the table.
      - **IPv4 Address or IPv6 Address:** Select this to add an IPv4 or IPv6 address entry. Click the **Value** field and enter the address. The default permission is **Allow**, which means that the appliance allows access to and from this device. You can change this to **Deny** to block access.
      - **IPv4 Network or IPv6 Network:** Select this to add an IPv4 or IPv6 network entry. Click the **Value** field and enter the network. The default permission is **Allow**, which means that the appliance allows syslog messages sent by this network. You can change this to **Deny** to block access.
      - **Any Address/Network:** Select this to allow or deny access to all IPv4 and IPv6 addresses and networks. The default permission is **Allow**, which means that the appliance allows syslog messages sent by all addresses and networks. You can change this to **Deny** to block access.

After you have added access control entries, you can perform the following:

- Select the ACEs that you want to group and put into a named ACL. Click the Create new named ACL icon and enter a name in the *Convert to Named ACL* dialog box.
  - Reorder the list of ACEs using the up and down arrows next to the table.
  - Select an IPv4 network and click the Edit icon to modify the entry.
  - Select an ACE and click the Delete icon to delete the entry. You can select multiple ACEs for deletion.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Setting DNS Logging Categories

You can specify logging categories you want the syslog to capture. Furthermore, you can filter these messages by severity at the Grid and member levels. For information about severity types, see [Specifying Syslog Servers](#).

To specify logging categories, complete the following:

1. From the **Data Management** tab, select the **DNS** tab, and then click **Grid DNS Properties** from the Toolbar.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click **Toggle Expert Mode** if the editor is in the basic mode, select the **Logging** tab, and then complete the following:
  - **Logging Facility:** Select a facility from the drop-down list. This is the location on the syslog server to which you want to sort the DNS logging messages.
  - **Logging Category:** Select one or more of these log categories:
    - **general:** Records the BIND messages that are not specifically classified.
    - **client:** Enables the logging of messages related to query processing, but not the queries themselves. Examples of messages include exceeding recursive client quota, and other errors related to recursive clients, blacklist and NXDOMAIN interception, query name rewrite, and others.
    - **config:** Records the configuration file parsing messages.
    - **database:** Records BIND's internal database processes.
    - **dnssec:** Records the DNSSEC-signed responses.
    - **lame servers:** Records bad delegation instances.
    - **network:** Records the network operation messages.
    - **notify:** Records the asynchronous zone change notification messages.
    - **queries:** Records the DNS queries. Note that enabling the logging of queries and responses will significantly affect system performance. Ensure that your system has sufficient CPU capacity before you enable DNS query logging.
    - **rate-limit:** Logs RRL (Response Rate Limiting) events. You must enable RRL in order for the appliance to log RRL events to this logging category.
    - **resolver:** Logs messages related to outgoing queries from the 'named' process, when it is acting as a resolver on behalf of clients.
    - **responses:** Records DNS responses. Note that enabling the logging of queries and responses will significantly affect system performance. Ensure that your system has sufficient CPU capacity before you enable DNS response logging.
    - **rpz:** Records log messages when responses are modified through RPZs or for which explicit passthru were invoked in the RPZs. This checkbox is not selected by default.
    - **security:** Logs miscellaneous messages that are related to security, such as denial or approval (mostly denial) of certain operations.
    - **transfer-in:** Records zone transfer messages from the remote name servers to the appliance.
    - **transfer-out:** Records zone transfer messages from the NIOS appliance to remote name servers.
    - **update:** Records the dynamic update instances.
    - **update-security:** Records the security updates.
    - **DTC load balancing:** Records information about which client is directed to which server.
    - **DTC health monitors:** Records any changes to the health state of a monitored server.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Viewing the Syslog

1. From the **Administration** tab, select the **Logs** tab -> **Syslog** tab.
2. From the drop-down list at the upper right-hand corner, select the Grid member on which you want to view the syslog.
3. Optionally, use the filters to narrow down the system messages you want to view. Click **Show Filters** to enable the filters. Configure the filter criteria, and then click **Apply**.  
Based on your filter criteria (if any), Grid Manager displays the following in the *Syslog* viewer:

•

: The Action icon column is displayed only when you have installed the RPZ license. Click this to view threat details in the *RPZ Threat Details* dialog box. For more information, see [Viewing the RPZ Threat Details](#) below.

- **Timestamp:** The date, time, and time zone of the log message. The time zone is the time zone configured on the member.
- **Facility:** The location on the syslog server that determines the processes and daemons from which the log messages are generated.
- **Level:** The severity of the message. This can be ALERT, CRITICAL, DEBUG, EMERGENCY, ERROR, INFO, NOTICE, or WARNING.
- **Server:** The name of the server that logs this message, plus the process ID.
- **Message:** Detailed information about the task performed. For Cloud Network Automation, this contains comma separated values of the admin, source, action, object, object type, and message values. Note that source is defined only if the cloud API request was proxied by the Cloud Platform Appliance. The format for this field is `proxied from:host,IP` where `host` and `IP` are the host name and IP address of the proxy.



#### Note

If the selected member is an HA pair, the Grid Manager displays the syslog in two tabs — **Active** and **Passive**.

Click the corresponding tab to view the syslog for each node.

## Viewing the RPZ Threat Details

Make sure that DNS resolution is enabled and running properly on the member to view threat details. To view threat details for the RPZ zones being queried, complete the following:

1. From the **Administration** tab, select the **Logs** tab -> **Syslog** tab.
2. Click the Action icon and select **View Threat Context** to open the *RPZ Threat Details* dialog. The **View Threat Context** option is disabled if there is no RPZ rule.
  - **RPZ Rule:** Displays the name of the RPZ rule.
  - **First Identified:** The date and timestamp of the first occasion that the threat was detected.
  - **Short Description:** The brief description of the threat.
  - **Description:** The description of the RPZ rule.  
Note the *RPZ Threat Details* dialog box may display *Unknown* if the threat is unknown, or *Unavailable* if the threat is known and threat details are not available.
3. Click the Close icon to close the *RPZ Threat Details* dialog.

You can also perform the following in the *Syslog* viewer:

- Toggle between the single line view and the multi-line view for display.
- Navigate to the next or last page of the file using the paging buttons.
- Refresh the syslog output with newly logged messages.
- Click the Follow icon to have the appliance automatically refresh the log every five seconds.
- Clear the contents of the syslog.
- Use filters and the **Go To** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go To** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria.
- To filter Microsoft synchronization related events, click **Show Filter**, select **Server** from the first drop-down list, and select **MS\_Server** from the drop-down list in the value field. This filter displays entries that begin with the prefix **ms**. To view values that belong to a specific Microsoft server, you must specify either the name or IP address of a given Microsoft server in the **Message** field. When you filter the syslog for a specific Grid member, it displays the log entries of Microsoft servers that are assigned to the respective Grid member when the entries are logged.
- Print the report or export it in CSV format.
- Bookmark the syslog page.

## Searching in the Syslog

Instead of paging through the syslog to locate messages, you can have the appliance search for syslog messages with certain text strings. To search for specific messages, perform the following:

- Enter a search value in the search field below the filters, and then click the **Search** icon. The appliance searches through the syslog and highlights the search value in the viewer. You can use the arrow keys next to the Search icon to locate the previous or next message that contains the search value.

## Downloading the Syslog File

To download the syslog file to a specified directory, if you want to analyze it later, complete the following:

1. From the **Administration** tab, select the **Logs** tab -> **Syslog** tab, and then click the Download icon.
2. Navigate to a directory where you want to save the file, optionally change the file name (the default names are *node\_1\_sysLog.tar.gz* and *node\_2\_sysLog.tar.gz*), and then click **OK**. If you want to download multiple syslog files to the same location, rename each downloaded file before downloading the next.



### Note

If your browser has a pop-up blocker enabled, you must turn off the pop-up blocker or configure your browser to allow pop-ups for downloading files.

## Enabling Automated Traffic Capture

Sometimes, there may be a sudden and unexplained change in the KPIs (Key Performance Indicators) in your network that may degrade your network performance and traffic capture during these events may not be available. For example, the cache hit ratio may drop below the configured value, or there may be an increase in authoritative delay.

You can configure NIOS so that a traffic capture may be triggered based on thresholds configured for DNS Cache Hit Ratio and Queries Per Seconds. You can then analyse the traffic capture data and use it to gather production data thus bringing down the time taken for root cause analysis. You can also attach the traffic capture data to a support case so that Infoblox Support can take the investigation forward.

You can choose to receive an SNMP trap or an email notification every time traffic capture is enabled or disabled or a support bundle is downloaded. For more information see, [Configuring SNMP](#).



### Note

If DNS Cache Acceleration is enabled on the IB-FLEX platform, the cache hit ratio is not updated for a minute thus triggering a false traffic capture.

To configure automated traffic capture:

1. From the **Grid** tab, select the **Grid Manager** tab -> click **Grid Properties** -> **Edit** from the Toolbar. Or to configure automated traffic capture for a member, from the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> select the member -> click the Action icon and click **Edit**.
2. In *Grid Properties Editor (for Grid)* or *Grid Member Properties Editor* (for members), click **Monitoring** -> **Advanced** tab.
3. Select the **Enable Automated Traffic Capture** checkbox.
4. In the **Capture Duration** field, enter the number of seconds for which traffic must be captured. For example, if you enter 30 in the **Capture Duration** field, traffic is captured for 30 seconds and stored in the traffic capture file. After 30 seconds, traffic is captured once again and the details are stored in a new traffic capture file. A new traffic capture is started only after the earlier traffic capture file is packaged. You can set a value from 1 to 600 seconds. **Capture Duration** is a mandatory field. If the capture file exceeds 2 GB (for IB-40XX and PT-40XX appliances) and 1 GB (for all other appliances), NIOS stops the traffic capture.
5. Select the **Save Local Copy** checkbox to save traffic capture data in your local disk.
6. Select the **Include Support Bundle** checkbox to download the support bundle after the traffic capture is complete and the system has remained in a stable state for 5 minutes.
7. From the **Export to** drop-down list, select your local management system (FTP server or SCP server) to which to transfer the traffic capture file. The default is **None**.
8. In the **Directory Path For Capture** field, enter the destination directory on the NIOS server for which to transfer the traffic capture files.

9. In the **Directory Path For Support Bundle** field, enter the destination directory on the NIOS server for which to transfer the support bundles.  
Note that if you enable the **Include Support Bundle** checkbox but do not specify the directory path for the support bundle, NIOS transfers the support bundles to the directory path that you specify in the **Directory Path For Capture** field.
10. In the **Server Address** field, enter the IP address of the FTP or the SCP server that you selected from the **Export to** drop-down list.
11. In the **Username** and **Password** fields, enter the user credentials to upload to the FTP or SCP server.
12. Select the **Enable Cache Hit Ratio Trigger** checkbox if you want NIOS to trigger a traffic capture based on the value of the cache hit ratio of recursive queries. If the cache utilization goes above the value you specify in the **Cache Utilization** field and the cache hit ratio goes below the value you specify in the **Hit Ratio Threshold Trigger** field, traffic capture is triggered. Once the cache hit ratio reaches the value you specify in the **Reset** field, NIOS stops the traffic capture.
13. Select the **Enable Queries Per Second Trigger** checkbox if you want NIOS to monitor the QPS (queries per second). NIOS triggers a traffic capture if the QPS value goes below the threshold value you specify in the **QPS Threshold Trigger** field. Once the QPS value reaches the value you specify in the **Reset** field, NIOS stops the traffic capture. For information about QPS, see [About Dashboards](#).
14. Select the **Enable Outgoing Recursive Queries Trigger** checkbox if you want NIOS to monitor the count of concurrent outgoing recursive queries. NIOS triggers a traffic capture if the count goes above the value you specify in the **Recursive Queries Threshold Trigger** field. If the number of concurrent recursive queries goes below the value you specify in the **Reset** field, NIOS stops the traffic capture. For information about recursive queries, see [Enabling Recursive Queries](#).
15. Select the **Enable Authoritative DNS Latency Trigger** checkbox if you want NIOS to monitor the authoritative DNS latency. NIOS performs the DNS latency on the IPS address you choose from the **Query IP Address** field and triggers a traffic capture if the DNS latency goes above the value you specify in the **Authoritative DNS Latency Threshold Trigger** field. The DNS latency is determined by querying a reverse zone against the IP address you selected using the `dig` command. If the authoritative DNS query latency goes below the value you specify in the **Reset** field, NIOS stops the traffic capture. For information about DNS latency, enabling and disabling DNS alert monitoring, see [Monitoring Tools](#).
16. Select the **Enable Recursive DNS Latency Trigger** checkbox if you want NIOS to monitor the recursive DNS latency. NIOS performs the DNS latency on the IPS address you choose from the **Query IP Address** field and triggers a traffic capture if the DNS latency goes above the value you specify in the **Recursive DNS Latency Threshold Trigger** field. The DNS latency is determined by querying each domain against the IP address you selected using the `dig` command. If the recursive DNS query latency goes below the value you specify in the **Reset** field, NIOS stops the traffic capture. For information about DNS latency, enabling and disabling DNS alert monitoring, see [Monitoring Tools](#).
17. Click **Save & Close**.

## Tracking Object Changes in the Database

If you have external applications that use information in the NIOS database, you can use the Object Change Tracking feature to get informed about changes in the NIOS database. You can then periodically synchronize IPAM, DNS, and DHCP data through the Infoblox API or RESTful API, which returns updated object information. When you enable this feature, the appliance tracks the changes that are made to NIOS objects. It assigns sequence IDs to all the changed objects. These sequence IDs are incremented when there is a change in the high-level objects such as IPv4 and IPv6 fixed addresses, networks, network containers, and others. When you query using the `db_objects` through the Infoblox API for desired object types, the application returns all the objects of those object types that changed after the sequence ID given in the query.

When you enable the Object Change Tracking feature, you can also specify the lifetime of deleted objects in the database and the total number of objects that must be present in the database. The deleted objects that are saved in the database are purged periodically. Note that you cannot track the user who created or updated an existing object. Infoblox supports full and incremental synchronization for these changes. Certain events such as Grid Master Candidate promotions and upgrades require a full synchronization.

Users with read-only permission can use this feature. For more information, see [Administrative Permissions for Object Change Tracking](#).

To enable and use the Object Change Tracking feature to track and synchronize updates, complete the following:

1. Enable the feature, as described in [Enabling Object Change Tracking](#) below.
2. Select whether you want to use a full synchronization or an incremental synchronization. When you use one of these synchronization methods to synchronize data through Infoblox API or RESTful API, they return updated objects that can be used to update your relational database. For full synchronization, see [Using Full Synchronization](#) below. For incremental synchronization, see [Using Incremental Synchronization](#) below.

## Best Practices for Object Change Tracking

- The Object Change Tracking feature is optimized to reduce impact on the DDI services and it runs only on the Grid Master. The synchronization process synchronizes 1000 objects at a time with a 2 second pause in between. There might be a slight impact on the Grid Master Candidate as they get updates from the Grid Master. When protocol services are running on the Grid Master Candidate you might encounter a 5% drop in the protocol performance. This feature does not impact the services that are running on the Grid members.
- NIOS does not update the sequence ID of the respective parent when you modify a child object. For example, inserting an A record under DNS zone does not increase the sequence number of the zone object.
- When you delete a parent object, NIOS saves the child objects in the deleted object table. But, if the associated child objects are either resource records or leases, NIOS will neither save them in the deleted object table nor update the sequence ID for these child objects.
- If you update a parent object, then the sequence ID of its child objects will not change.
- NIOS updates the sequence ID of the host record and IPv4 and IPv6 host addresses, if there are any changes to host addresses, both IPv4 and IPv6. The sequence ID is also updated when you update the host alias records that results in increase in the sequence ID of the host record and all its child objects.

For example, if you create a host record `hhh.test.com` with an IPv4 address and perform an incremental synchronization to get updates, the response contains an updated host record and host address:

```
curl -k1 -u admin:infoblox -H content-type:application/json -X GET
https://10.32.2.202/wapi/v2.5/db_objects?
start_sequence_id=2830689052:0\&object_types=
record:host,record:host_ipv4addr\;_return_fields=last_sequence_id,object,obj
ect_type\;
_return_type=json-pretty
```

The response is as follows:

```
[
  {
    "_ref": "db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDIw:38",
    "object": {
      "_ref": "record:host/
ZG5zLmhvc3QkLl9kZWZhdWx0LmNvbS5rYXJqYWdpLmhvc3Qx:
hhh.test.com/default",
      "ipv4addrs": [
        {
          "_ref":
"record:host_ipv4addr/
ZG5zLmhvc3RfYWRRkcmVzcyQuX2RlZmF1bHQuY29tLmthcmphZ2kuaG9zdDEuMS4x
LjEuMS4:1.1.1.1/ hhh.test.com/default",
```

```

        "configure_for_dhcp": false,
        "host": " hhh.test.com",
        "ipv4addr": "1.1.1.1",
        "mac": "11:11:11:11:11:11"
    }
],
    "ipv6addrs": [
        {
            "_ref":
"record:host_ipv6addr/
ZG5zLmhvc3RfYWRRkcmVzcyQuX2RlZmF1bHQuY29tLmthcmphZ2kuaG9zdDEuYWE6
OmFhLg:aa%3A%3Aaa/ hhh.test.com/default",
            "configure_for_dhcp": false,
            "host": " hhh.test.com",
            "ipv6addr": "aa::aa"
        }
    ],
    "name": " hhh.test.com",
    "view": "default"
},
    "object_type": "record:host",
    "unique_id": "c326fcf8058c4022939050af96a0fdb2"
},
{
    "_ref": "db_objects/Li5hbGxfY2hhbmdlZmF9vYmplY3RzJDIx:38",
    "object": {
        "_ref":
"record:host_ipv6addr/
ZG5zLmhvc3RfYWRRkcmVzcyQuX2RlZmF1bHQuY29tLmthcmphZ2kuaG9zdDEuYWE6
OmFhLg:aa%3A%3Aaa/ hhh.test.com/default",
        "configure_for_dhcp": false,
        "host": " hhh.test.com",
        "ipv6addr": "aa::aa"
    },
    "object_type": "record:host_ipv6addr",
    "unique_id": "2b19a399c3f747538d879fdd33a4de32"
},

```



```

{
  "_ref": "db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDIy:38",
  "object": {
    "_ref":
"record:host_ipv4addr/
ZG5zLmhvc3RfYWwRkcmVzcyQuX2RlZmF1bHQuY29tLmthcmphZ2kuaG9zdDEuMS4x
LjEuMS4:1.1.1.1/ hhh.test.com/default",
    "configure_for_dhcp": false,
    "host": " hhh.test.com",
    "ipv4addr": "1.1.1.1",
    "mac": "11:11:11:11:11:11"
  },
  "object_type": "record:host_ipv4addr",
  "unique_id": "3fb8fba003d647ceac5796a28459725a"
},
{
  "_ref": "db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDIz:3956302770%3A38",
  "last_sequence_id": "3956302770:38"
}
]

```

- When you update DNS host, host address, and host alias, NIOS updates the sequence IDs of child objects associated with these parent objects even though the changes do not affect the child objects. For example, when you update an existing comment in the host record, NIOS updates the sequence IDs of child objects associated with the respective host record.

Example:

```

curl -k1 -u admin:infoblox -H content-type:application/json -X GET
https://10.32.2.202/wapi/v2.5/db_objects?
start_sequence_id=2830689052:0&object_types=
record:host,record:host_ipv4addr\;_return_fields=last_sequence_id,object,obj
ect_type,o
bject.comment\;_return_type=json-pretty

```

The response is as follows:

```

[
  {
    "_ref": "db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDI:27",
    "object": {
      "_ref":
"record:host/ZG5zLmhvc3QkLl9kZWZhdWx0LmNvbS50ZXN0LmhoaA:hhh.test.com/

```

```

default",
  "comment": "hi"
},
  "object_type": "record:host"
  "unique_id": "1ab6c989d8fd454ca06ffac6ac600fe6"
},
{
  "_ref": "db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDM:27",
  "object": {
    "_ref":
"record:host_ipv4addr/
ZG5zLmhc3RfYWRRkcmVzcyQuX2RlZmF1bHQuY29tLnRlc3QuaGhoLjEuMC4wLjEu
:1.0.0.1/hhh.test.com/default",
    "configure_for_dhcp": false,
    "host": "hhh.test.com",
    "ipv4addr": "1.0.0.1"
  },
  "object_type": "record:host_ipv4addr"
  "unique_id": "1ab6c989d8fd454ca06ffac6ac600fe7"
},
{
  "_ref": "db_objects/
Li5hbGxfY2hhbmdlZF9vYmplY3RzJDQ:2830689052%3A27",
  "last_sequence_id": "2830689052:27"
}

```

- When an IPv4 or an IPv6 host address changes, NIOS deletes the old record and creates a new record with the updated IP address. Note that new sequence IDs are generated for the associated objects. When you query a host record, NIOS displays the list of host addresses associated with it.

Example:

```

curl -k1 -u admin:infoblox -H content-type:application/json -X GET
https://10.32.2.202/wapi/v2.5/db_objects?
start_sequence_id=2830689052:0&object_types=
record:host,record:host_ipv4addr\;_return_fields=last_sequence_id,object,obj
ect_type\;
_return_type=json-pretty

```

The response is as follows:

```

[
  {
    "_ref": "db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDA:28",
    "object": {
      "_ref":
"record:host/ZG5zLmhvc3QkLl9kZWZhdWx0LmNvbS50ZXN0LmhoaA:hhh.test.com/
default",
      "ipv4addrs": [
        {
          "_ref":
"record:host_ipv4addr/
ZG5zLmhvc3RfYWwRkcmVzcyQuX2RlZmF1bHQuY29tLnRlc3QuaGhoLjEuMC4wLjIu
:1.0.0.2/hhh.test.com/default",
          "configure_for_dhcp": false,
          "host": "hhh.test.com",
          "ipv4addr": "1.0.0.2"
        }
      ],
      "name": "hhh.test.com",
      "view": "default"
    },
    "object_type": "record:host"
    "unique_id": "1ab6c989d8fd454ca06ffac6ac600fe6"
  },
  {
    "_ref": "db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDE:28",
    "object": {
      "_ref": "deleted_objects/Li5kZWxldGVkX29iamVjdHMkMA",
      "object_type": "record:host_ipv4addr",
      "unique_id": "68cad3406a244aa0b34b3ea3be383598"
    },
    "object_type": "deleted_objects"
    "unique_id": "1ab6c989d8fd454ca06ffac6ac600fe4"
  },
  {
    "_ref": "db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDI:28",
    "object": {

```

```

    "_ref":
    "record:host_ipv4addr/
ZG5zLmhvc3RfYWRRkcmVzcyQuX2RLZmF1bHQuY29tLnRlc3QuaGhoLjEuMC4wLjIu
:1.0.0.2/hhh.test.com/default",
    "configure_for_dhcp": false,
    "host": "hhh.test.com",
    "ipv4addr": "1.0.0.2"
  },
  "object_type": "record:host_ipv4addr"
  "unique_id": "1ab6c989d8fd454ca06ffac6ac600fe9"
},
{
  "_ref": "db_objects/
Li5hbGxfY2hhbmdlZF9vYmpLY3RzJDM:2830689052%3A28",
  "last_sequence_id": "2830689052:28"
}

```

- NIOS returns the shared records per zone when you query for them. For example, consider two zones, z1 and z2, and a shared record group, srg1, which contains a shared record sr1. The shared record group srg1 is associated with the zones z1 and z2. Hence, the shared record sr1 is shared between two zones, z1 and z2, but NIOS saves only a single copy of the shared record sr1 in the database. When you query for shared record updates, NIOS returns two records for z1 and z2, as sr1 is associated with both the zones, z1 and z2, and saves these records with a UUID and a sequence ID in the table. The response also contains new records with zone pointing to the zone it belongs to, UUID and sequence ID of the new records and a `shared_record_group`, which indicates that it is a shared record and not a normal resource record. When you enable Object Change Tracking feature, NIOS clears this table and recreates a new table for the existing shared records.
- Note that NIOS generates two leases in the database when a client sends a DHCP request to the NIOS DHCP server, which is configured with failover association. You can perform either a full or an incremental synchronization in this scenario. Infoblox recommends that you ignore the lease that is generated by the failover DHCP server with `binding_state` as `BACKUP` and consider the lease with `binding_state` as `ACTIVE`.
- File distribution service automatically removes old files from `wapi_output` directory if there is no sufficient space for new files. NIOS allocates 50% each of the file distribution area to `wapi_output` directory and normal files. For example, if the file distribution area is allocated 500 MB, then 250 MB is used for `wapi_output` directory and the remaining 250 MB is used for normal files.
- NIOS does not support WAPI paging mechanism. As the synchronization results are displayed in the order of `sequence_id`, you must specify `_max_results` to achieve paging mechanism for synchronization.
- NIOS deletes all previous scheduled tasks after a restore or an upgrade operation.

## Enabling Object Change Tracking

To enable object change tracking:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Grid Properties** -> **Edit** from the Toolbar.

**Standalone appliance:** From the **System** tab → **System Manager** tab, expand the Toolbar and click **System Properties** -> **Edit**.

2. In the *Grid Properties* editor, click **Object Change Tracking** and complete the following:
  - **Enable Object Change Tracking:** Select this checkbox to enable object change tracking for the objects.
  - **Status:** Indicates whether object change tracking is enabled or not. You can view what percentage of the operation is complete.

Note this operation might take a longer time to complete when you have a large database. For example, 1 million objects on an IB-4015 appliance takes approximately 1 hour. Infoblox recommends that you enable this feature during off-peak hours.

- **Maximum number of deleted objects that will be tracked:** Specify the total number of deleted objects that must be present in the database. The minimum value is 2,000 and the maximum value is 20,000. The default value is 4,000.
3. Save the configuration.

## Using Full Synchronization

Full synchronization synchronizes the entire set of objects irrespective of updates. You can perform a full synchronization only on the Grid Master and Grid Master Candidate. Note that the tasks created on the Grid Master and Grid Master Candidates will be executed on them respectively and you can download the synchronized data on to the member where the task was created. The Grid Master Candidate can execute only read-only tasks. To save the output of a full synchronization, you must specify the output location. The output of a full synchronization is in the following formats: JSON, XML, ROWJSON, or ROWXML.

You can specify one of the following output locations: `"FILE_DISTRIBUTION"` and `"LOCAL"`.

- When `_output_location = "FILE_DISTRIBUTION"`

When you set the output location as mentioned above, you must specify a file name or a prefix for the file. NIOS saves the output file in the file distribution area and displays an error message if you do not specify a file name or a prefix. For more information about file distribution area, see [Managing Directories](#).

When you select the above output location for a Grid Master Candidate, you must select **Allow Upload to Grid Members** to upload files to the file distribution area. Infoblox recommends that you use the Grid Master Candidate to offload the Grid Master. You must enable a full synchronization during off-peak hours if protocol services are running on the Grid Master Candidate. For more information, see [Enabling Upload to Grid Members](#).

Example:

```
curl -k1 -u admin:infoblox -X POST 'https://.../wapi/v2.5/fileop?
_function=read' -H 'Content-Type: application/json; charset=UTF-8' -d
'{"_filename": "test.json",
 "_output_location": "FILE_DISTRIBUTION", "_object": "db_objects",
 "all_object_types_supported_in_version": "2.5", "_encoding": "ROWJSON"}'
```

You can specify either `_object = db_objects` or `all_object_types_supported_in_version=2.5` to retrieve all objects. When you specify `_object = db_objects`, NIOS returns only standard Restful API fields as response. To retrieve all fields of all objects, you can specify

`all_object_types_supported_in_version=2.5`. To fetch the output file, start file distribution service and use the following URL:

```
curl -k1 -u admin:infoblox -X GET http://<ipaddress>/wapi_output/test.json -H
'Content-Type:application/json'
```

- When `_output_location = "LOCAL"`
  - If you schedule a full synchronization and set the output location as mentioned above, a file name is generated based on the task ID and the output file is saved at another location. You do not have to specify a file name or a prefix.
  - If you do not schedule a synchronization, then NIOS returns an URL and a token, which is used for the download complete function, in the response.
  - If you do not specify an output location, NIOS saves the file in the file distribution area.
  - Use the above output location for saving full synchronization files when file distribution area does not have enough space for storage.

Example:

```
curl -k1 -u admin:infoblox -X POST 'https://127.0.0.1/wapi/v2.5/fileop?
_function=read&_schedinfo.schedule_now=true' -H 'Content-Type: application/
json; charset=UTF-8' -d '{"_output_location": "LOCAL", "_object":
"db_objects", "all_object_types_supported_in_version": "2.5", "_encoding":
"JSON"}'
```

NIOS saves the output file in the file distribution area if you do not specify an output location as shown in the code below:

```
curl -k1 -u admin:infoblox -X POST 'https://.../wapi/v2.5/fileop?
_function=read' -H 'Content-Type: application/json; charset=UTF-8' -d
'{"_filename": "test.json",
"_output_location": "FILE_DISTRIBUTION", "_object": "db_objects",
"all_object_types_supported_in_version": "2.5", "_encoding": "ROWJSON"}'
```

Note the following about full synchronization:

- A full synchronization is required initially for a complete snapshot of the database.
- Depending on the size of requested data, a full synchronization may take a longer time to complete. For example, when 1 million objects are requested with all fields on an IB-4015 appliance, full synchronization takes a couple of hours. When 1 million objects are requested with standard RESTful API fields on an IB-4015 appliance, full synchronization takes less than an hour.
- Infoblox recommends that you request only object types with standard RESTful API fields and specify any additional required fields during a full synchronization.
- To dump full synchronization updates into a file using `fileop->read` operation when the updated objects are large in number, you can specify `_object=db_objects` and all the necessary parameters.
- Infoblox recommends a full synchronization in the following cases:
  - After an upgrade, restore or master promotion, which resets the sequence ID.
  - During off-peak hours for busy Grids with high rate of change that is greater than 500/sec.
  - When the maximum time or maximum number to track deleted objects is met, with the default being 4 hours and 4000 objects, then synchronization API returns an error if the last deleted sequence ID is newer than the current sequence ID.

Following are a few samples of API requests for full synchronization:

- If you specify `_output_location = "local"` and schedule the task, then you must use `fileop ->get_file_url` to retrieve the URL:

```
curl -k1 -u admin:infoblox -X POST 'https://127.0.0.1/wapi/v2.5/fileop?
_function=get_file_url' -H "Content-Type: application/json" -d
'{"task_id":1}'
```

The response is as follows:

```
"url":  
"https://10.32.2.202/http_direct_file_io/req_id-OBJECTS-1/nios-db-objects-1.json"
```

You can fetch this file using the following API request:

```
curl -k1 -u admin:infoblox -X GET https://10.32.2.202/http_direct_file_io/req_id-OBJECTS-1/nios-db-objects-1.JSON -H 'Content-Type: application/json; charset=UTF-8'
```

- If you specify `_output_location = "local"` and do not schedule the task, then NIOS returns an URL and a token:

```
curl -k1 -u admin:infoblox -X POST 'https://127.0.0.1/wapi/v2.5/fileop?_function=read'  
-H 'Content-Type: application/json; charset=UTF-8' -d  
{ "_output_location": "LOCAL", "_object": "db_objects",  
"all_object_types_supported_in_version": "2.5", "_encoding": "JSON" }
```

The response is as follows:

```
{  
  "token":  
    "eJy1kU1vwyAMhu/8kfaSD/JB0t5aZZU2Ta3UTtrRSob2n1JggUzstv5+ZtJ123QFk/  
L7mMUZK6+4w\n6QujTVr  
jwzTLYCfmOFtKNGc7jPaWWqPjCnenPev60MNRn5krmAQYZhwDGgCmUAbmSrZUrmKnhb45\nnn04Q8  
KoXzNVsxyt RtKVYFU0qqqpuBwf+tJinkWRBBW8hOL/OMp6nZZ2KtG2ymAKF1FuAM44a0GaT\n/  
gBUSXd43T8fNl3C85xnBq1 P1AB2eCevT5J8leR1UuRcQFGsy3pN1KfTYU+sJmJRudQS9a/  
rSFpF\nnk6KnUsxz8mWe5tJfdBau7n/64vyHCdp  
Iq9BcYrYg+Pbx21D+Gq5Wxany00hu87KB48Munmvmw9FxnET+BNyTi0DtA4+YAn3ryaE20tWzvh  
/QLT4GbhW=  
=\n",  
  "url":  
    "https://10.35.6.87/http_direct_file_io/req_id-DOWNLOAD-1001/nios-  
db_objects--09-05-20  
16_22:35:27.JSON"  
}
```

- To save the file in a local area instead of the file distribution area:

```
curl -k1 -u admin:infoblox -X POST 'https://127.0.0.1/wapi/v2.5/fileop?_function=read&_schedinfo.schedule_now=true' -H 'Content-Type: application/
```



```
json; charset=UTF-8' -d '{"_output_location": "LOCAL", "_object":
"db_objects", "all_object_types_supported_in_version": "2.5", "_encoding":
"JSON"}'
```

The response is as follows:

```
scheduledtask/b25\LnF1ZXV\ZF90YXNrJDE:1/WAITING_EXECUTION
```

- You can also schedule a full synchronization task as follows:

```
curl -k1 -u admin:infoblox -X POST 'https://127.0.0.1/wapi/v2.5/fileop?
_function=read&_schedinfo.schedule_now=true' -H 'Content-Type: application/
json; charset=UTF-8' -d '{"_filename": "test2.txt", "_object": "db_objects",
"all_object_types_supported_in_version": "2.5", "_encoding": "JSON"}'
scheduledtask/b25\LnF1ZXV\ZF90YXNrJDE:1/WAITING_EXECUTION
```

NIOS returns the scheduled task reference so that you can query and find out when will the scheduled task be complete:

```
curl -k1 -u admin:infoblox -X GET
```

```
https://127.0.0.1/wapi/v2.5/scheduledtask/b25\LnF1ZXV\ZF90YXNrJDE:1/
WAITING_EXECUTION
```

The response is as follows:

```
{ "_ref": "scheduledtask/b25\LnF1ZXV\ZF90YXNrJDE:1/COMPLETED",
"approval_status": "NONE", "execution_status": "COMPLETED", "task_id": 1 }
```

Note that the result is in either JSON or XML format. You can fetch the output file from the `wapi-output` directory of the file distribution area when the scheduled task is complete.

## Using Incremental Synchronization

Incremental synchronization synchronizes only those objects that are updated since the previous synchronization. With incremental synchronization, you can query NIOS for any updates and when updates are found, NIOS returns the changed objects since the previous incremental synchronization. Incremental synchronization synchronizes the NIOS database for object changes such as addition, updates, or deletion, mostly via scripts. The incremental synchronization query consists of object types, a cookie of the form

`<id of db:sequence_id>`, a sequence ID and an optional `exclude_deleted` object to indicate if the deleted objects must be excluded from the results. NIOS includes the deleted records in the results by default.

### Note

Infoblox recommends that you use `fileop->read` for incremental synchronization when the updated objects are large in number.

Note the following about incremental synchronization:

- You cannot request an incremental synchronization to get updated (created, updated or deleted) objects after a full synchronization.

- For performance reasons, NIOS returns the objects that changed and not the object data that changed. Infoblox recommends that you figure out what changed by comparing the object with the previous version in the client database.
- Incremental synchronization may take a longer time to complete depending on the number of changes. For example, if 50K objects are requested with all fields on an IB-4015 appliance, an incremental synchronization approximately takes 2 minutes. When 50K objects are requested with standard RESTful API fields on an IB-4015, the synchronization process completes in approximately 1 minute.
- The size of the memory increases when the number of objects increases. Infoblox recommends that you use `_max_results` to limit the number of results or use the asynchronous `fileop->read` operation.
- Infoblox recommends that for a busy Grid with a high Rate of Change that is greater than 500/sec, you must run an incremental synchronization more frequently. Example: If the expected Rate of Change for objects in the Grid is 50/sec, you can schedule an incremental synchronization every 5-10 minutes. If the Rate of Change for objects is 200/sec, schedule an incremental synchronization every minute or every 2 minutes.
- Infoblox recommends that you include all object types when you query using an incremental synchronization, so that it returns the highest sequence ID for all objects that can be used for subsequent incremental updates. If you query only selected object types such as A record and zone, NIOS returns the highest sequence ID for those object types only.  
Consider an example where A record, Network, and A record are the object types with sequence ID 4, 5 and 6 respectively. When you query for an A record object type with a previous sequence ID using incremental synchronization, NIOS returns A record with the highest sequence ID, which is 6 and does not return the network object. Meanwhile, if you add or update a network object, then in the subsequent incremental synchronization it does not return the network object and you will lose updates made to the network object. Hence, Infoblox recommends that you include all object types in the query to ensure that you do not lose updates made to other object types.
- You cannot perform a full or an incremental synchronization during Grid Master Candidate promotions, restores, and upgrades.
- The output of an incremental synchronization is either in JSON or XML format.

The following are a few samples of API requests for incremental synchronization:

- To get updates with `all_objects_supported_in_version=2.5`:

```
curl -k1 -u admin:infoblox -H content-type:application/json -X GET
https://10.35.6.87/wapi/v2.5/db_objects?
start_sequence_id=992684967:0\&all_object_type
s_supported_in_version=2.5\;
return_fields=last_sequence_id,object,object_type,unique
_id\;_return_type=json
```

The response for the above API request contains all fields of all object types.

- To get an A record and an auth zone object update:

```
curl -k1 -u admin:infoblox -H content-type:application/json -X GET
https://10.32.2.202/wapi/v2.5/db_objects?
start_sequence_id=183566281:0\&object_types=r
ecord:a,zone_auth\;_return_fields=last_sequence_id,object,object_type,unique
_id\;_return_type=json
```

The response is as follows:

```
[{"_ref": "db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDA:19", "object":
{"_ref":
```

```
"zone_auth/ZG5zLnpvbmUkLl9kZWZhdWx0LmNvbS5mb28:foo.com/default", "fqdn":
"foo.com",
"view": "default"}, "object_type": "zone_auth", "unique_id":
"087af628fa03418faa0577b953807efc"}, {"_ref":
"db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDE:21", "object": {"_ref":
"record:a/
ZG5zLmJpbmRfYSQuX2RlZmF1bHQuY29tLmZvbyxhcmVjMSwxLjIuMy4y:arec1.foo.com/
default", "ipv4addr": "1.2.3.2", "name": "arec1.foo.com", "view":
"default"}, "object_type":
"record:a", "unique_id": "89314c3fae2841f49600e1b686b0c7b7"}, {"_ref":
"db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDI:183566281%3A21",
"last_sequence_id":
"183566281:21"}},
```

- NIOS returns deleted objects as synthetic deleted objects. To get updates on deleted host objects:

```
curl -s -k1 -u admin:infoblox -H content-type:application/json -X GET
https://10.34.19.220/wapi/v2.5/db_objects?
start_sequence_id=1207501969:0\&object_types
=record:host,bulkhost,record:host_ipv4addr,record:host_ipv6addr\&exclude_dele
ted=false
```

The response is as follows:

```
[{ "_ref": "db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDM2YQ:130",
  "object": {
    "_ref": "deleted_objects/Li5kZWxldGVkX29iamVjdHMkMA",
    "object_type": "record:host_ipv6addr"
  },
  "object_type": "deleted_objects",
  "unique_id": "f7d87ec4b9204e16a3d14bfb340c402d" },
{ "_ref": "db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDM2Yg:130",
  "object": {
    "_ref": "deleted_objects/Li5kZWxldGVkX29iamVjdHMkMQ",
    "object_type": "record:host_ipv4addr" },
  "object_type": "deleted_objects",
  "unique_id": "68316853600f41f083e2be134628ce62" },
{ "_ref": "db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDM2Yw:130",
  "object": {
    "_ref": "deleted_objects/Li5kZWxldGVkX29iamVjdHMkMg",
```

```

    "object_type": "record:host" },
    "object_type": "deleted_objects",
    "unique_id": "7a5e0622e9db441caa47c12ca30700d5" },
{ "_ref": "db_objects/Li5hbGxfY2hhbmdlZF9vYmplY3RzJDM2ZA:1207501969%3A130",
  "last_sequence_id": "1207501969:130" }]

```

- To dump incremental synchronization updates into a file using `fileop->read` operation when the updated objects are large in number, you can specify `_object=db_objects` and all the necessary parameters for incremental synchronization:

```

curl -k1 -u admin:infoblox -X POST 'https://10.34.19.100/wapi/v2.5/fileop?
_function=read' -H 'Content-Type: application/json; charset=UTF-8' -d
'{"_output_location": "LOCAL", "_object": "db_objects",
"start_sequence_id" :
"1973553522:0","all_object_types_supported_in_version": "2.5",
"_max_results": 1000000,"_encoding": "JSON"}'

```

For more information about supported objects in the Infoblox API and Restful API, refer to the Infoblox API Documentation and the Infoblox WAPI Documentation.

## Configuring the PTop Instance to Collect CPU Utilization Data

The NIOS appliance allows you to configure the collection of CPU utilization data of top processes that are running on the Grid Master and all members of a Grid, or a standalone system. The PTop tool that runs in the background, collects the CPU utilization data and logs it in PTop logs.

To configure the PTop instance, complete the following steps:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab. Expand the Toolbar, and then select **Grid Properties -> Edit**.  
Or,  
**Standalone system:** From the **System** tab, select the **System Manager** tab. Expand the Toolbar, and then select **System Properties -> Edit**.
2. In the *Grid Properties* editor or the *System Properties* editor, click the **Monitoring** tab, and then complete the following steps in the *PTop Log* section:
  - **Number of Top Processes:** Enter the number of top processes for which you want to check the CPU utilization. If you enter 11, NIOS displays the CPU utilization of the top 11 processes. The minimum number of processes for which you can check the utilization is 10 and the maximum is 25. By default, the widget displays the CPU utilization for 10 top processes.
  - **PTop Interval:** Enter the time interval in seconds to determine the frequency with which the Ptop tool must run and collect data. The Ptop tool is an internal tool that runs in the background and collects CPU utilization and other data. The minimum interval that you can enter is 2 seconds and the maximum is 30 seconds. By default, the PTop interval is set as 2 seconds.
3. Click **Save & Close**.

## Collecting Database Performance Data

The NIOS appliance provides Ptop logs that consist of database metrics that you can use to determine the health of the NIOS database and baseline its performance. The metrics help you to ascertain the impact of changes such as adding a Grid member or enabling features such as Grid replication for DNS zones or multi-master DNS, on the database performance.

You can download the Ptop logs by using the WAPI calls and then use the logs for further analysis. These WAPI calls allow you to download only the Ptop logs instead of the whole support bundle.

To download the Ptop log files, perform the following steps. The steps are showing sample WAPI requests.

1. Run the following command to get the URL for downloading the Ptop logs:

```
$curl -k1 -u admin:infoblox -H "Content-Type: application/json" -X POST https://10.12.21.107/wapi/v2.9.7/fileop?_function=get_log_files -d '{"include_rotated":false,"log_type":"PTOPLLOG","node_type":"ACTIVE"}'
```

The output includes the token for this transaction and the URL from where to download the Ptop logs as follows:

```
{
  "token":
  "eJytUEFuwyAQvPOR5BLM2jEkvaVyI1WqEimp10PKBuwi2YYCqZK+vhA2jZ2WFmR0rrbuj1\nQN
  I17Ryiv8hoPXFAltLMve1Ge6V21vnEm90BNG1s8aR74koiEbuLGa0ZEYkyMhJXkaVya3Je6Ksz\
  n/
  obRTHpBXE32UIs156LiQDeMiRK2JJwXFz8mmCfCe4wuPBQFMAo\loyVQYKLIXVQm2YvYmj+HV60
  u2YFUNaMV5yB4ExUdeGidaMdaGw9Hb7S/yJLGZWi5/G7t5UjXVwBKnCCmX\ndtBFnNw/
  mQL4EUEv0zw97f07JiGeHPCCoogE9kK03kWn9nbcgh79+1Ds3sLhE/tQ/
  GzhnbJqy6567j3wFfoFk=\n",
  "url": "https://10.12.21.107/http_direct_file_io/req_id-
  DOWNLOAD-1125063601760735/ptoplog.tar.gz"
}
```

2. Download the Ptop log files using the obtained URL:

```
$curl -k1 -u admin:infoblox -H "Content-type:application/force-download" -O https://10.12.21.107/http_direct_file_io/req_id-
DOWNLOAD-1125063601760735/ptoplog.tar.gz
```

3. To ensure that the URL does not remain open for a long time, use the token obtained in step 1 and close the URL as follows:

```
$curl -k -u admin:infoblox -H 'Content-Type:application/json' -X POST https://10.12.21.107/wapi/v2.9.7/fileop?_function=downloadcomplete -d '{"token":
  "eJytUEFuwyAQvPOR5BLM2jEkvaVyI1WqEimp10PKBuwi2YYCqZK+vhA2jZ2WFmR0rrbuj1\nQN
  I17Ryiv8hoPXFAltLMve1Ge6V21vnEm90BNG1s8aR74koiEbuLGa0ZEYkyMhJXkaVya3Je6Ksz\
  n/
  obRTHpBXE32UIs156LiQDeMiRK2JJwXFz8mmCfCe4wuPBQFMAo\loyVQYKLIXVQm2YvYmj+HV60
  u2YFUNaMV5yB4ExUdeGidaMdaGw9Hb7S/yJLGZWi5/G7t5UjXVwBKnCCmX\ndtBFnNw/
  mQL4EUEv0zw97f07JiGeHPCCoogE9kK03kWn9nbcgh79+1Ds3sLhE/tQ/
  GzhnbJqy6567j3wFfoFk=\n"}'
```

## Configuring dnstap

You can use the dnstap log format to log DNS queries and responses at high rates to well-known destinations. NIOS logs all valid DNS queries and responses that are not dropped by Advanced DNS Protection. For information about dnstap, see <https://dnstap.info/>.

For information about capturing DNS queries and responses without using dnstap, see [Capturing DNS Queries and Responses](#).



### Note

dnstap for high-performance query logging is supported on Infoblox appliances on which Advanced DNS Protection or DNS Cache Acceleration is running.

This section includes the following topics:

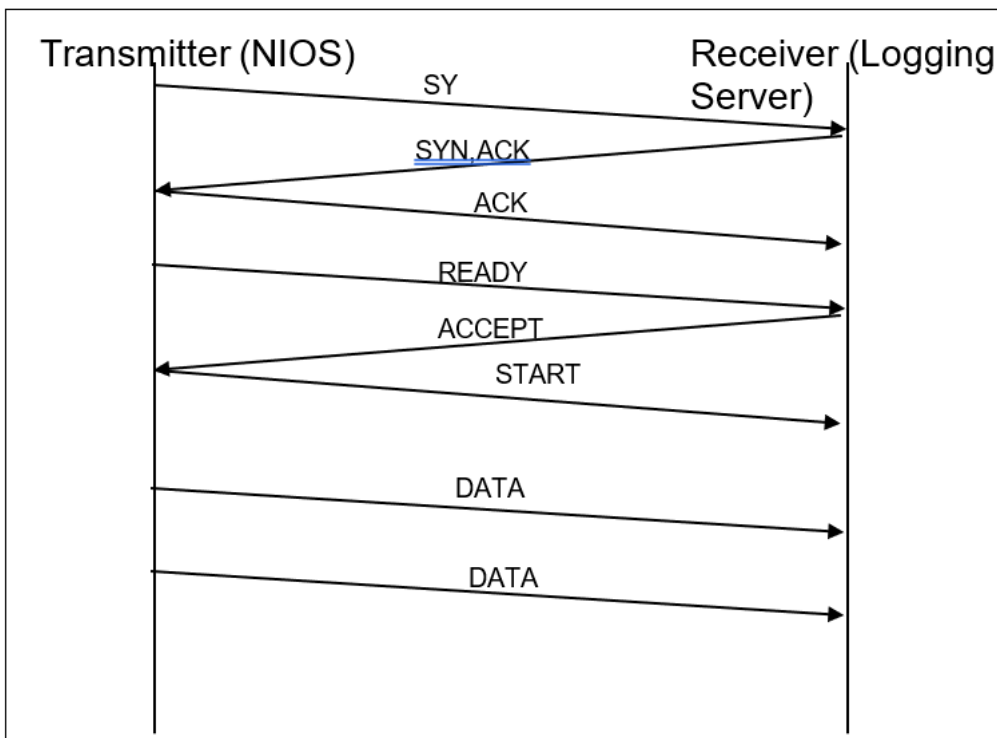
- [About dnstap Implementation](#)
- [Protobuf Template for Infoblox dnstap](#)
- [Configuring dnstap to Log DNS Queries and Responses](#)

### About dnstap Implementation

dnstap is implemented using the Google protocol buffer (protobuf) for packaging logged data, and Frame Streams as a lightweight streaming protocol to transmit the data to a receiver.

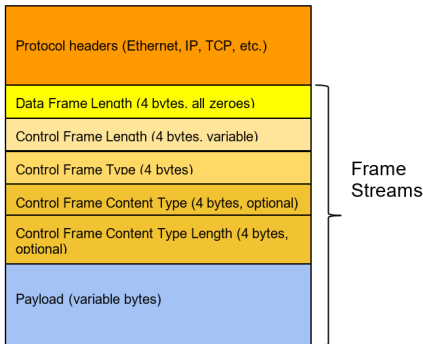
A bi-directional communication starts with a 3-way handshake sequence of Frame Streams control packets over a TCP (or other reliable byte-streaming). The following figure illustrates the connection between the transmitter and receiver pair.

*Packet Flow for Frame Streams 3-way handshake initialization and data frames over a TCP connection*



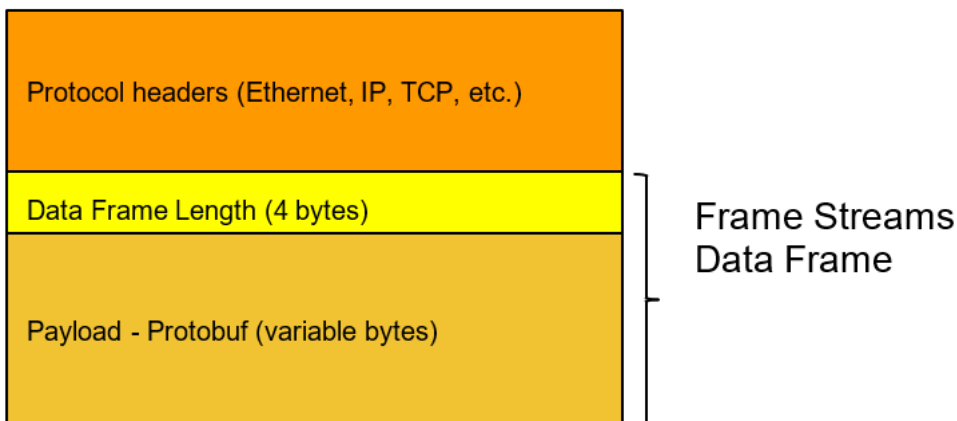
The transmitter initiates the communications to the receiver by first sending a READY control frame carrying the `protobuf:dnstap.Dnstap` string. The receiver responds with an ACCEPT control frame carrying the same `protobuf:dnstap.Dnstap` string. The handshake is completed when the transmitter sends a START control frame carrying the same `protobuf:dnstap.Dnstap` string.

*Frame Streams Control Frame Format*



Data packets are carried over the connection with only a 4-byte framing overhead as illustrated in the following figure.

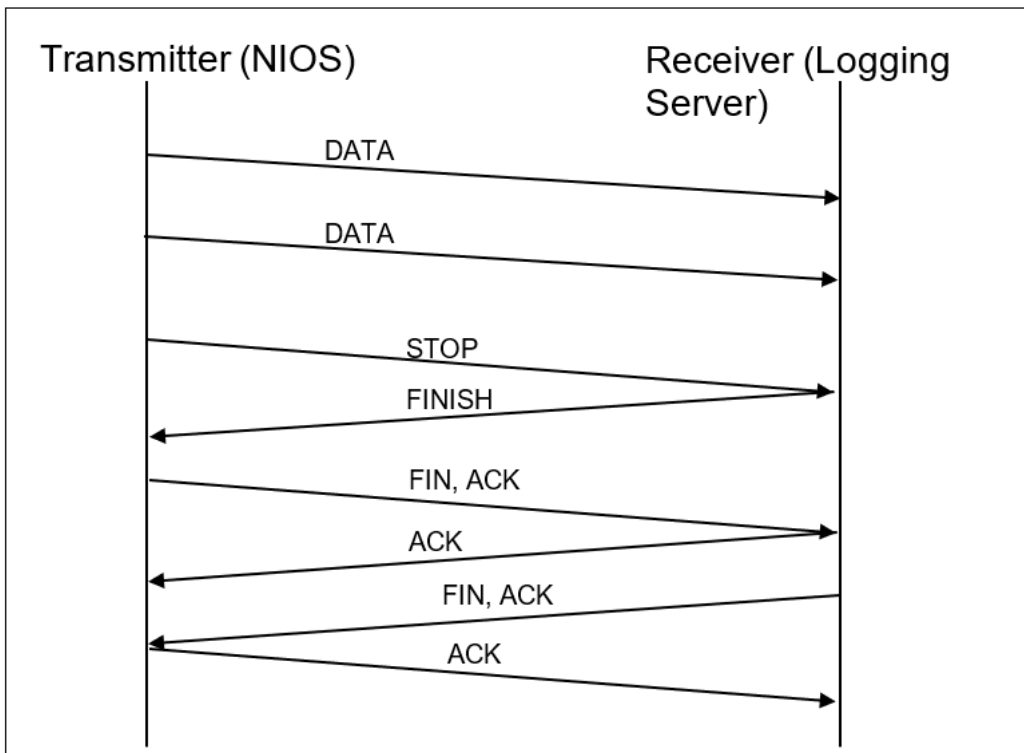
*Frame Streams Data Frame Format*



To end the bi-directional communication that is shown in the following figure, the sequence starts with a STOP control frame from one of the communicating pairs, followed by a FINISH acknowledgment control frame from the other half of the pair.

*Packet flow for Frame Streams termination over a TCP connection*





#### Protobuf Template for Infoblox dnstap

The following is a sample template that you can create for flexible and efficient use of your logged data. You can specify the way in which information must be logged by defining protobuf message types.

```
// dnstap: flexible, structured event replication format for DNS software,
// This file contains the protobuf schemas for the "dnstap" structured event
// replication format for DNS software.
// Written in 2013-2014 by Farsight Security, Inc.
// Updated in 2019-2020 by Infoblox, Inc.
// To the extent possible under law, the author(s) have dedicated all copyright
// and related and neighboring rights to this file to the public domain
worldwide.
// This file is distributed without any warranty.
// You should have received a copy of the CC0 Public Domain Dedication along
with this file.
// If not, see:
// < http://creativecommons.org/publicdomain/zero/1.0/ >. package dnstap;
// "Dnstap": this is the top-level dnstap type, which is a "union" type that
// contains other kinds of dnstap payloads, although currently only one type
// of dnstap payload is defined.
```

```

// See: https://developers.google.com/protocol-buffers/docs/techniques#union
message Dnstap {
  // DNS server identity.
  optional bytes identity = 1;
  // DNS server version.
  optional bytes version = 2;
  // Extra data for this payload.
  optional bytes extra = 3;
  // Identifies which field below is filled in.
  enum Type {
    MESSAGE = 1;
  }
  required Type type = 15;
  // One of the following will be filled in.
  optional Message message = 14;
}
// SocketFamily: the network protocol family of a socket. This specifies how
// to interpret "network address" fields.
enum SocketFamily {
  INET = 1; // IPv4 (RFC 791)
  INET6 = 2; // IPv6 (RFC 2460)
}
// SocketProtocol: the transport protocol of a socket. This specifies how to
// interpret "transport port" fields.
enum SocketProtocol {
  UDP = 1; // User Datagram Protocol (RFC 768)
  TCP = 2; // Transmission Control Protocol (RFC 793)
}
// Message: a wire-format (RFC 1035 section 4) DNS message and associated
// metadata. Applications generating "Message" payloads should follow
// certain requirements based on the MessageType, see below.
message Message {
  // We are supporting the following types of messages
  // CQ: CLIENT_QUERY
  // CR: CLIENT_RESPONSE
  enum Type {
    // CLIENT_QUERY is a DNS query message sent from a client to a DNS

```

```

// server which is expected to perform further recursion, from the
// perspective of the DNS server. The client may be a stub resolver or
// forwarder or some other type of software which typically sets the RD
// (recursion desired) bit when querying the DNS server. The DNS server
// may be a simple forwarding proxy or it may be a full recursive
// resolver.
CLIENT_QUERY = 5;
// CLIENT_RESPONSE is a DNS response message sent from a DNS server to
// a client, from the perspective of the DNS server. The DNS server
// typically sets the RA (recursion available) bit when responding.
CLIENT_RESPONSE = 6;
}
// One of the Type values described above.
required Type type = 1;
// One of the SocketFamily values described above.
optional SocketFamily socket_family = 2;
// One of the SocketProtocol values described above.
optional SocketProtocol socket_protocol = 3;
// The network address of the message initiator.
// For SocketFamily INET, this field is 4 octets (IPv4 address).
// For SocketFamily INET6, this field is 16 octets (IPv6 address).
optional bytes query_address = 4;
// The network address of the message responder.
// For SocketFamily INET, this field is 4 octets (IPv4 address).
// For SocketFamily INET6, this field is 16 octets (IPv6 address).
optional bytes response_address = 5;
// The transport port of the message initiator.
// This is a 16-bit UDP or TCP port number, depending on SocketProtocol.
optional uint32 query_port = 6;
// The transport port of the message responder.
// This is a 16-bit UDP or TCP port number, depending on SocketProtocol.
optional uint32 response_port = 7;
// The time at which the DNS query message was sent or received, depending
// on whether this is an AUTH_QUERY, RESOLVER_QUERY, or CLIENT_QUERY.
// This is the number of seconds since the UNIX epoch.
optional uint64 query_time_sec = 8;
// The time at which the DNS query message was sent or received.

```

```

// This is the seconds fraction, expressed as a count of nanoseconds.
optional fixed32 query_time_nsec = 9;
// The initiator's original wire-format DNS query message, verbatim.
optional bytes query_message = 10;
// This is a wire-format DNS domain name.
// Currently, we are not supporting this.
optional bytes query_zone = 11;
// The time at which the DNS response message was sent or received,
// depending on whether this is an CLIENT_RESPONSE.
// This is the number of seconds since the UNIX epoch.
optional uint64 response_time_sec = 12;
// The time at which the DNS response message was sent or received.
// This is the seconds fraction, expressed as a count of nanoseconds.
optional fixed32 response_time_nsec = 13;
// The responder's original wire-format DNS response message, verbatim.
optional bytes response_message = 14;
// Start of Infoblox specific log messages
optional bytes subscriber_id = 15; // 32 bytes
optional uint64 local_id = 16; // 8 bytes
optional bytes pcp_ssp = 17; // 16 bytes
optional bytes proxy_all = 18; // Proxy all
optional bytes fqdn = 19; // MAX 255 bytes
required uint32 txid = 20; //16-bit transaction id
}
// All fields except for 'type' in the Message schema are optional.
// It is recommended that at least the following fields be filled in for
// particular types of Messages.
// CLIENT_QUERY:
// socket_family, socket_protocol
// query_message
// query_time_sec, query_time_nsec
// CLIENT_RESPONSE:
// socket_family, socket_protocol
// query_time_sec, query_time_nsec
// response_message
// response_time_sec, response_time_nsec

```

## Configuring dnstap to Log DNS Queries and Responses

You can use the dnstap log format to log DNS queries and responses at high rates to well-known destinations and achieve high performance query and response logging. NIOS logs all valid DNS queries and responses that are not dropped by Advanced DNS Protection. For information about dnstap, see <https://dnstap.info/>.

To use the dnstap log format, you need to enable dnstap by running the `set enable_dnstap` command. To view whether the dnstap log format is enabled or disabled, run the `show dnstap-status` command. To view the number of queries and responses sent to the destination when the dnstap log format is enabled for high-performance logging of queries and responses, run the `show dnstap-stats` command.

If you choose to enable the dnstap log format, you will not be able to capture queries and responses using the **Data connector for all DNS Queries/Responses to a Domain** fields. And if you use the **Data connector for all DNS Queries/Responses to a Domain** fields for query capture, the **DNSTAP settings for DNS Queries/Responses** fields will be disabled.



### Note

For Advanced DNS Protection software with acceleration, you must download the latest ruleset before enabling dnstap.

## Limitations of Using dnstap to Log Queries and Responses

Ensure that you understand the following limitations before you use dnstap to log queries and responses:

- dnstap supports UDP, TCP, and EDNS protocols that require additional processing thus leading to a decrease in performance.
- NIOS does not support BIND9 dnstap.
- If the remote logging server is not accessible, then the logs are dropped and not buffered.
- The dnstap server cannot truncate EDNS0 queries.
- If you run a query that contains `+edns=1`, a dnstap server that uses the Golang DNS library to process the captured data displays it as a bad signature (TSIG signature failure).
- Capturing the queries and responses also depends on other factors such as the size of the flavor deployed and features enabled over it.
- dnstap does not support query and response logging on the MGMT interface.

## Configuring dnstap to Log DNS Queries and Response Captures

To configure dnstap to log DNS queries and to capture responses, complete the following:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the **Toolbar**, and then click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab, and then click the **Members** tab -> *member* checkbox -> Edit icon.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click **Toggle Advanced Mode** and then select the **Logging** tab.
3. Select the **Queries** checkbox to start capturing DNS queries. When you enable this option at the member level, NIOS captures DNS queries for only the selected member.
4. Select the **Responses** checkbox to start capturing DNS responses. When you enable this option at the member level, NIOS captures DNS responses for only the selected member.
5. In the **DNSTAP Receiver Address** field, enter the IP address from which you want to capture queries or responses. It supports both IPv4 and IPv6 addresses.
6. In the **DNSTAP Receiver Port** field, enter the port number on which you want to configure the dnstap client system. The default port number is 6000.
7. Click **Save and Close**.

Infoblox recommends the configurations in the following table to meet high-performance query logging using the dnstap log format:

Feature	Total CPU	Total Virtual Memory (without Advanced DNS Protection software)	Total Virtual Memory (with Advanced DNS Protection software)	Database Object Count	Grid Master Capable
Small recursive DNS (with acceleration)	10	16	24	100,000	No
Medium recursive DNS (with acceleration)	16	24	32	100,000	No
Large recursive DNS (with acceleration)	26	34	42	100,000	No

## Configuring DHCP Fingerprints and Fingerprint Detection

This section explains the Infoblox DHCP fingerprint detection feature and how to configure it on the appliance. It also explains how to configure DHCP fingerprints for IPv4 and IPv6. It contains the following sections:

- [DHCP Fingerprint Detection](#)
- [Enabling and Disabling DHCP Fingerprint Detection](#)
- [Standard and Custom DHCP Fingerprints](#)
- [Configuring DHCP Fingerprints](#)

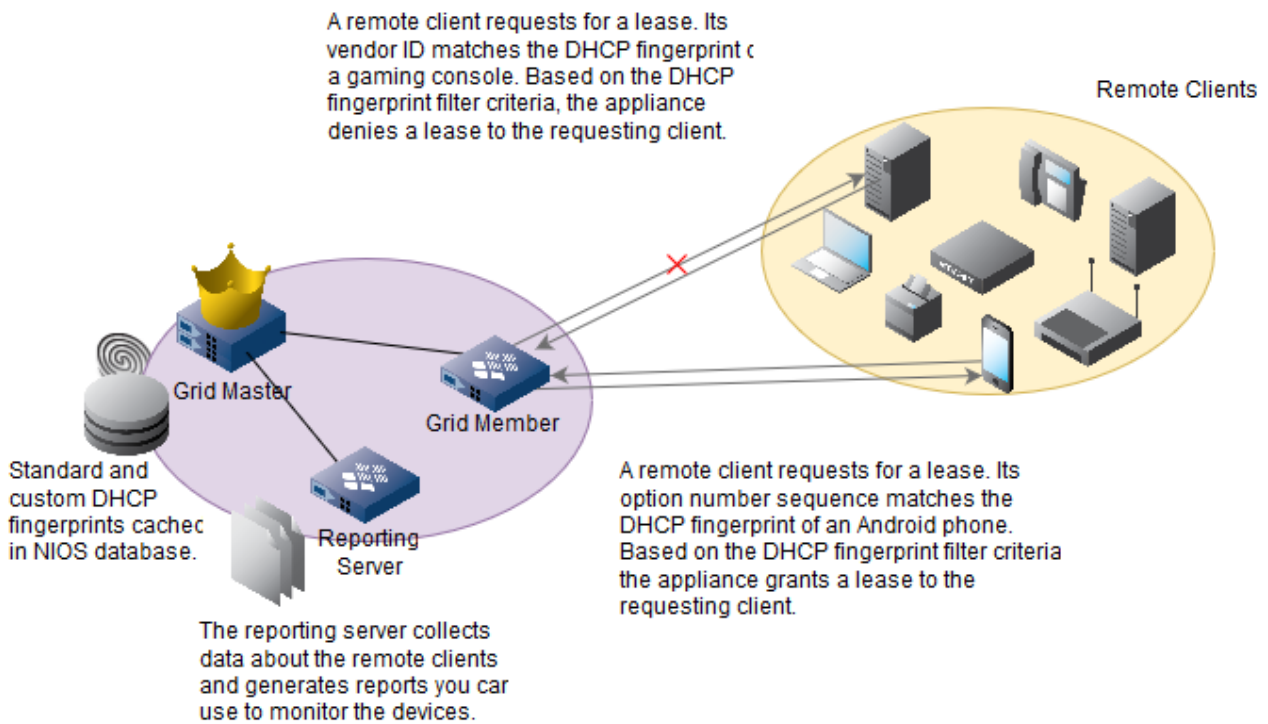
### DHCP Fingerprint Detection

The NIOS appliance utilizes DHCP fingerprint detection to identify IPv4 and IPv6 mobile devices such as laptop computers, tablets and smart phones, on your network. Due to the broadcast and pervasive nature of DHCP, using DHCP fingerprint detection is an efficient way to perform system identification and inventory. You can use DHCP fingerprint detection to track devices on your network, block those that are not allowed (such as gaming consoles and home routers), and plan for future growth by accessing trending information such as the number of Apple iPhones versus that of Android phones.

When a remote DHCP client sends a DHCP REQUEST message, it includes a set of DHCP options, such as option 55 and 60. Option 55 contains an option number sequence the appliance uses to interpret the list of DHCP options that the client requests. The appliance returns the values of these requested options if the information is available. Option 60 contains a value that indicates the device type of the requesting client. Information in option 55 or 60 is incorporated to form a unique identifier known as the DHCP fingerprint, which the appliance uses to identify the requesting client.

On an Infoblox appliance, DHCP fingerprint detection is enabled by default for all new installations. You can disable this feature at the Grid and member levels. For information, see [Enabling and Disabling DHCP Fingerprint Detection](#). As illustrated in the below figure, the appliance automatically matches option 55 and then option 60 in DHCP REQUEST messages against standard and custom DHCP fingerprints in the database. Once the appliance finds a match, it either grants or denies a lease to the requesting client based on the DHCP fingerprint filters that you apply to the DHCP range. For information about how to configure DHCP fingerprints, see [Configuring DHCP Fingerprint Filters](#). For information about how to define and apply DHCP fingerprint filters, see [Configuring DHCP Fingerprint Filters](#). To obtain trending information about the top OSs (operating systems) or vendor IDs for remote clients, Infoblox provides a few reports from which you can extract data. For information about reports, see [Infoblox Reporting and Analytics](#).

## DHCP Fingerprint Detection



## About DHCP Fingerprints

When a DHCP client sends a REQUEST message and includes DHCP option 55 (the parameter request list) and option 60 (the vendor identifier), it provides information about its OS and the device type. The combination of the option sequence or vendor ID in option 55 or 60 is used to infer the OS and device type of the remote client. These parameters are then incorporated into a DHCP fingerprint that provides unique information about this client.

For example, the option number sequence for a Microsoft Windows Kernel 5.1 and 5.2 systems in option 55 can be one of the following:

- 1, 15, 3, 6, 44, 46, 4
- 1, 15, 3, 6, 44, 46, 47, 31, 33, 121, 249, 43, 200
- 1, 15, 3, 6, 44, 46, 47, 31, 33, 249, 43
- 1, 15, 3, 6, 44, 46, 47, 31, 33, 249, 43, 252
- 1, 15, 3, 6, 44, 46, 47, 31, 33, 249, 43, 252, 12
- 15, 3, 6, 44, 46, 47, 31, 33, 249, 43
- 15, 3, 6, 44, 46, 47, 31, 33, 249, 43, 252
- 15, 3, 6, 44, 46, 47, 31, 33, 249, 43, 252, 12
- 28, 2, 3, 15, 6, 12, 44, 47

The option number sequence for an Apple Mobile Device can be one of the following:

- 1, 121, 3, 6, 15, 119, 252, 0, 0, 0, 0, 0, 0, 0, 0, 0
- 1, 121, 3, 6, 15, 119, 252, 118
- 1, 3, 6, 12, 15, 28, 44, 121



1, 3, 6, 15, 19

1, 3, 6, 21, 1, 25, 2, 82

Additionally, DHCP option 60 tracks the vendor ID. The vendor ID can be generic or specific. For example, the vendor ID `MSFT 5.0` for a Microsoft Windows Kernel 5.1 or 5.2 system and a Microsoft Windows Kernel 6.0 system can be the same. For certain Microsoft Windows Kernel 5.0 systems, the vendor ID can be `MSFT`, which is generic, or it can be `MSFT 5.0`, which is specific. Depending on how specific the option number sequence and the vendor ID are, this information can form a unique identifier, the DHCP fingerprint for a remote client.

#### Note

If you have enabled a firewall and if the corresponding firewall rules or policies are set to modify options 55 and 60 of the remote DHCP client to mask the identity of the client, then fingerprinting clients cannot be achieved.

## Enabling and Disabling DHCP Fingerprint Detection

Grid DHCP fingerprint detection is enabled by default for new installations, and no special licenses are required. You can disable this or override the Grid setting at a member level. Note that when you enable DHCP fingerprint detection, there will be a slight impact on DHCP performance.

When you enable DHCP fingerprint on an HA pair, both peers in a failover association maintain the same DHCP fingerprinting state (enabled or disabled) even when one of the peers fails or becomes operational again. Note that both peers must be in the same Grid for the fingerprinting state to stay the same. For information about DHCP failover, see [DHCP Failover](#).

To enable and disable Grid DHCP fingerprinting:

1. **Grid:** From the **Data Management** tab, select the **DHCP** tab, and then click **Grid DHCP Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DHCP** tab -> **Members** tab -> **Members** -> *member* checkbox, and then click the Edit icon.
2. In the *Grid DHCP Properties* or *Member DHCP Properties* editor, select the **Fingerprinting** tab.
3. Complete the following
  - **Enable Fingerprint Detection:** Deselect this checkbox to disable the feature. You can enable DHCP fingerprint detection again by selecting the checkbox. Click **Override** to change the configuration for a member, or click **Inherit** to inherit the Grid setting.
4. Save the configuration and click **Restart** at the top of the screen.

## Standard and Custom DHCP Fingerprints

Standard DHCP fingerprints are automatically installed on the appliance when you first set it up or after you have completed an upgrade from previous NIOS releases to NIOS 6.7 and later. Note that new DHCP fingerprints are added to the appliance during a major NIOS upgrade. For more information about upgrades, see [Upgrading NIOS](#).

#### Note

When you upgrade to NIOS 6.7 and later, DHCP fingerprint detection is disabled during the upgrade. You must enable it if you want the appliance to use DHCP fingerprint detection. For information, see [Enabling and Disabling DHCP Fingerprint Detection](#).

You can configure custom DHCP fingerprints for devices whose DHCP fingerprints are not captured in the standard DHCP fingerprints. For information about how to add custom DHCP fingerprints, see as described in [Adding New DHCP Fingerprints](#).

Both standard and custom DHCP fingerprints are cached in memory for matching purposes. Depending on the information provided in a DHCP fingerprint, the appliance first matches the option number sequence sent in the DHCP REQUEST message. If option 55 is not included in the request or if there is no match from the cached DHCP

fingerprints, the appliance then tries to match the vendor ID in option 60. When there is an option number sequence match, the appliance displays the name of the DHCP fingerprint in Grid Manager. If there is no option number sequence match but there is a vendor ID match, the appliance displays the vendor ID. For information about how to view fingerprint information, see as described in [Viewing DHCP Fingerprint Information](#).

You can also create IPv4 and IPv6 DHCP fingerprint filters and then apply them as class filters to specific IPv4 and IPv6 DHCP ranges and range templates. For information about how to configure and use DHCP fingerprint filters, see [Configuring DHCP Fingerprint Filters](#).

## Administrative Permissions

DHCP fingerprint detection is enabled by default for new installations. For upgrades, you must enable this feature after the upgrade is completed. For information, see [Enabling and Disabling DHCP Fingerprint Detection](#). No special licenses are required for this feature.

Superusers can add, modify, and delete DHCP fingerprints and DHCP fingerprint filters. Limited-access users with Read/Write permission to DHCP fingerprints can add, modify, and delete DHCP fingerprints while those who have Read-only permission can only view information in the **Data Management** tab -> **DHCP** tab -> **Fingerprints** tab. For information about administrative permissions, see [About Administrative Permissions](#).

## Configuring DHCP Fingerprints

The appliance installs standard DHCP fingerprints when you first install or upgrade to NIOS 6.7 or later. You cannot modify standard DHCP fingerprints nor delete them, but you can disable them. When you disable a DHCP fingerprint, the appliance disables the associated option number sequence and vendor ID, and it cannot match a remote device against the disabled DHCP fingerprint.

When you add a new DHCP fingerprint, the appliance marks it as a custom DHCP fingerprint. For information about adding custom DHCP fingerprints, see [Adding New DHCP Fingerprints](#) below. You can modify information about custom DHCP fingerprints, and you can delete them. For information, see [Modifying Custom DHCP Fingerprints](#) and [Deleting Custom DHCP Fingerprints](#) below. When you delete a custom DHCP fingerprint, the appliance moves it to the Recycle Bin, if enabled. You can later restore it from the Recycle Bin if needed.

Activities such as additions, modifications, and deletions of DHCP fingerprints are recorded in the audit log. For information about how to use the audit log, see [Monitoring Tools](#).

### Note

The appliance periodically updates the cached DHCP fingerprints. When you add, modify, or delete a DHCP fingerprint, you do not need to restart the services, but it may take up to two minutes before the appliance updates the DHCP fingerprint.

## Adding New DHCP Fingerprints

To add a custom DHCP fingerprint, complete the following steps:

1. From the **Data Management** tab, select the **DHCP** tab -> **Fingerprints** tab, and then click the Add icon.
2. In the *Add DHCP Fingerprint* wizard, complete the following:
  - **Name**: Enter the name of the custom DHCP fingerprint. The name must be unique, and it cannot contain any UTF-8 characters.
  - **Device Class**: From the drop-down list, select the device category to which this new fingerprint belongs. You can also enter a new device class here. When you enter a device class that already exists, the appliance matches the entry and uses the class from the current list. Device class is used for filtering purposes. For information about defining a DHCP fingerprint filter, see [Configuring DHCP Fingerprint Filters](#).
  - **Protocol**: From the drop-down list, select the protocol used for custom DHCP fingerprint.
  - **Option Number Sequence**: Click the Add icon in the table. The appliance adds a row to the table. Click the row and enter the DHCP option number that you want the appliance to validate. Valid values are from 0 to 255. When you enter more than one option, you must use commas (without spaces) to separate the numbers. For example, you can enter `1, 15, 3, 6, 44, 46, 47, 31, 33, 249, 43` for a Microsoft

Windows Kernel 5.1.5.2 system.

You can also select an option sequence and click the Delete icon to delete it. Note that if you enter an option sequence that already exists in a standard DHCP fingerprint, you must disable that standard fingerprint before you can add the option sequence to the new DHCP fingerprint.

- **Vendor Identifier:** Click the Add icon in the table. The appliance adds a row to the table. Click the row and enter a vendor ID for this fingerprint. You can add more than one vendor ID. You can also select a vendor ID and click the Delete icon to delete it.
- **Comment:** Enter additional information about the custom DHCP fingerprint.
- **Disabled:** Select this if you want to save the configuration for the DHCP fingerprint but do not want to activate it yet. You can clear this checkbox when you are ready to use this DHCP fingerprint.

3. Save the configuration or click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).

To schedule this task, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, click **Later**, and then specify a date, time, and time zone.

### Modifying Custom DHCP Fingerprints

You can modify custom DHCP fingerprints, but not the standard ones. Note that the appliance periodically updates the cached DHCP fingerprints. When you modify a DHCP fingerprint, you do not have to restart the services, but it may take up to two minutes before the appliance updates the DHCP fingerprint.

To modify a custom DHCP fingerprint, complete the following steps:

1. From the **Data Management** tab, select the **DHCP** tab -> **Fingerprints** tab -> *custom\_fingerprint* checkbox, and then click the Edit icon.
2. The *DHCP Fingerprint* editor provides the following tabs from which you can modify information:
  - **General:** Modify the general information, such as the name and device class, as described in Adding New DHCP Fingerprints above. Note that when you change the name of a DHCP fingerprint, the old name no longer exists, and you cannot use it for searching or filtering purposes. You may not be able to modify all fields in a standard DHCP fingerprint.
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with the DHCP fingerprint. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#). Note that you cannot modify extensible attributes for standard fingerprints.
3. Save the configuration.

To schedule this task, click the Schedule icon at the top of the wizard. In the *Schedule Change* panel, click **Later**, and then specify a date, time, and time zone.

### Deleting Custom DHCP Fingerprints

When you delete a custom DHCP fingerprint, the appliance moves it to the Recycle Bin, if enabled. You can later restore the DHCP fingerprint if needed. Note that you cannot delete standard DHCP fingerprints.

To delete a custom DHCP fingerprint, complete the following steps:

1. From the **Data Management** tab, select the **DHCP** tab -> **Fingerprints** tab -> *custom\_fingerprint* checkbox, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes** to delete the DHCP fingerprint.

To schedule this task, click the Delete icon -> **Schedule Delete**. In the *Schedule Deletion* dialog box, click **Delete Later**, and then specify a date, time, and time zone.

### Viewing DHCP Fingerprint Information

The following are a few ways that you can use to view the DHCP fingerprint information:

- In a Grid with a reporting server, you can view reports that contain information about the top OSs and device types of the leasing clients in your network. For more information about DHCP Dashboards, see [Predefined Dashboards](#).

- The appliance provides a few predefined smart folders from which you can view lease information about specific device groups, such as gaming consoles and Microsoft Windows devices. For more information, see [Predefined Smart Folders](#).
- When the appliance finds a DHCP fingerprint match for a client, Grid Manager displays either the fingerprint name or the vendor ID in the following panels of Grid Manager: *IP List*, *Current Lease*, *Lease History*, and *DHCP Range*. You can see this information in the **Fingerprints** column in these panels.
- The appliance records all DHCP fingerprint related activities in the audit log. For more information about using the audit log, see [Monitoring Tools](#).

## Monitoring with SNMP

This section describes how you can use SNMP (Simple Network Management Protocol) to monitor NIOS appliances in your network. It contains the following topics:

- [Understanding Simple Network Management Protocol](#)
- [Configuring SNMP](#)
- [SNMP MIB Hierarchy](#)

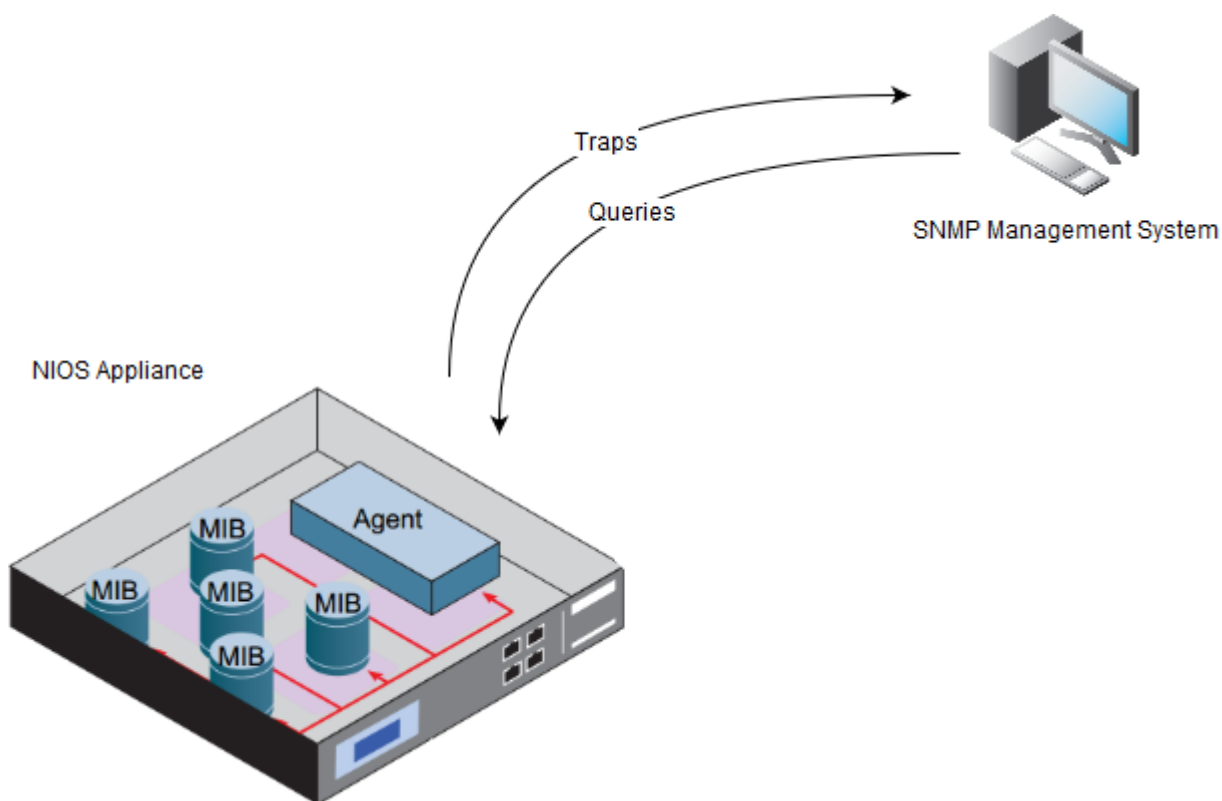
### Related topic

[Understanding Simple Network Management Protocol](#)

## Understanding Simple Network Management Protocol

You can use SNMP (Simple Network Management Protocol) to manage network devices and monitor their processes. An SNMP-managed device, such as a NIOS appliance, has an SNMP agent that collects data and stores them as objects in MIBs (Management Information Bases). The SNMP agent can also send traps (or notifications) to alert you when certain events occur within the appliance or on the network. You can view data in the SNMP MIBs and receive SNMP traps on a management system running an SNMP management application, such as HP OpenView, IBM Tivoli NetView, or any of the freely available or commercial SNMP management applications on the Internet.

*SNMP Overview*



The NIOS appliance supports SNMPv1, SNMPv2, and SNMPv3. It also adheres to the following RFCs:

- *RFC 3411, An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks*
- *RFC 3412, Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)*
- *RFC 3413, Simple Network Management Protocol (SNMP) Applications*
- *RFC 3414, User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMP)*
- *RFC 3416, Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)*
- *RFC 3418, Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*
- *RFC 1155, Structure and Identification of Management information for TCP/IP-based internets*
- *RFC 1213, Management Information Base for Network Management of TCP/IP-based internets:MIB-II*
- *RFC 2578, Structure of Management Information Version 2 (SMIv2)*

### About SNMPv1 and SNMPv2

SNMPv1 is the initial implementation of SNMP. It operates over protocols such as UDP (User Datagram Protocol) and IP (Internet Protocol). SNMPv2 includes improvements in performance and security. It adds new protocol operations such as GetBulk and Inform, which allow the management system to request larger blocks of data from the agent. Both SNMPv1 and SNMPv2 use common strings that are sent in clear text to authenticate clients.

The NIOS appliance supports SNMPv1 and SNMPv2 in which the SNMPv2 agent acts as a proxy agent for the SNMPv1 management systems. When an SNMPv1 management system sends a query to the appliance, the SNMPv2 proxy agent forwards the request to the SNMPv1 agent. The proxy agent maps the SNMPv1 trap messages to the SNMPv2 trap messages, and then forward the messages to the management system.

You can enable the appliance to receive queries from SNMPv1 and SNMPv2 management systems. You can also add SNMPv1 and SNMPv2 management systems to receive traps from the appliance. For information about how to configure SNMPv1 and SNMPv2 on the appliance, see [Configuring SNMP](#).

## About User-Based Security Model in SNMPv3

SNMPv3 adds security and remote configuration enhancements to SNMPv1 and SNMPv2. The NIOS appliance supports the USM (User-based Security Model) in SNMPv3 for the authentication, encryption, and decryption of SNMP data.

SNMPv3 uses the same MIB objects as those supported in SNMPv1 and SNMPv2.

SNMPv3 provides the following security measures:

- **Data integrity:** Ensure that SNMP data is not maliciously modified by unauthorized entities during its transmission through the network. This protects against unauthorized management operations, such as falsifying the value of a MIB object.
- **Authentication:** Verify the identities of the origin of the SNMP data to protect against masquerade threats that may temper the identity of users who have the appropriate authorization to send and receive SNMP data.
- **Confidentiality:** Ensure that unauthorized users cannot eavesdrop on any data exchanges between SNMP agents and management systems, depending on local policies of the systems.
- **Timeliness:** Ensure that the SNMP data is received in a timely manner to prevent malicious reordering of data by unauthorized entities.

To enable SNMPv3 on the NIOS appliance to provide user-based security, you must first configure SNMPv3 users on the appliance to enable access by SNMP management systems. The appliance supports HMAC-MD5-96 and HMAC-SHA-96 hash functions as the authentication protocols, and DES (Data Encryption Standard) and AES (Advanced Encryptions Standard) as the encryption methods for SNMPv3 users. For information, see [Configuring SNMP](#).

## Configuring SNMP

You can configure the appliance to receive SNMP queries from specific management systems and send SNMP traps to specific trap receivers. SNMP operation supports both IPv4 and IPv6 networks. The appliance supports SNMPv1, SNMPv2, and SNMPv3. You can set up either SNMPv1/SNMPv2 or SNMPv3, or all of them for a Grid. You can also override the Grid settings at a member level.

To configure SNMPv1 and SNMPv2 on the appliance, do the following:

- Enable the NIOS appliance to accept queries, as described in [Accepting Queries](#) in the following section.
- Specify the management systems to which the appliance sends traps, as described in [Adding Trap Receivers](#) in the following section.
- Specify system information using managed objects in MIB-II, the standard MIB defined in *RFC 1213*. For information, see [Setting SNMP System Information](#) in the following section.

To configure SNMPv3 on the appliance, do the following:

- Add an SNMPv3 user and set up authentication and privacy protocols. For information, see [Configuring SNMPv3 User](#) in the following section. After you set up an SNMPv3 user, you can modify and delete it.
- Enable the NIOS appliance to accept queries, as described in [Accepting Queries](#) in the following section.
- Specify the management systems to which the appliance sends traps, as described in see [Adding Trap Receivers](#) in the following section.
- Specify system information using managed objects in MIB-II, the standard MIB defined in *RFC 1213*.

## Configuring SNMPv3 Users

To enable SNMPv3, you must first configure SNMPv3 users on the appliance. For information about SNMPv3, see [Understanding Simple Network Management Protocol](#).

To configure an SNMPv3 user:

1. From the **Administration** tab, select the **SNMPv3 Users** tab, and then click the Add icon.
2. In the *Add SNMPv3 User* wizard, complete the following:
  - **Name:** Enter a user name for the SNMPv3 management system.
  - **Authentication Protocol:** Select one of the following:
    - **MD5:** Select this to use the HMAC-MD5-96 authentication protocol to authenticate the SNMPv3 user.  
This protocol uses the MD5 (Message-Digest algorithm 5) hash function in HMAC (Hash-based

Message Authentication Code) and truncates the output to 96 bits. The output is included as part of the SNMP message sent to the receiver. For detailed information about the protocol, refer to *RFC1321, The MD5 Message-Digest Algorithm*.

- **SHA:** Select this to use the HMAC-SHA-96 authentication protocol to authenticate the SNMPv3 user.  
This protocol uses the SHA (Secure Hash Algorithm) hash function and truncates the output to 96 bits. The output is included as part of the SNMP message sent to the receiver.
- **None:** Select this to decline using any authentication protocol for this SNMPv3 user. When you select this option, you are not required to enter a password.
  - **Password:** Enter a password for the selected authentication protocol.
  - **Confirm Password:** Enter the same password.
- **Privacy Protocol:** Select one of the following:
  - **DES:** Select this to use DES for data encryption. DES is a block cipher that employs a 56-bit key size and 64-bit block size in the encryption.
  - **AES:** Select this to use AES for data encryption. AES is a symmetric-key encryption standard that comprises AES-128 block cipher. The cipher has a 128-bit block size and a key size of 128 bits.
  - **None:** Select this to decline using any privacy protocol for this SNMPv3 user. When you select this option, you are not required to enter a password.
    - **Password:** Enter a password for the privacy protocol.
    - **Confirm Password:** Enter the same password.
- **Comment:** Enter useful information about the SNMP user, such as location or department.
- **Disable:** Select this checkbox to retain an inactive profile for this SNMP user in the configuration. You can clear this checkbox to activate the profile.

If an SNMPv3 user is configured to send SNMP queries, you cannot delete the user.

3. Click **Next** to define extensible attributes. For information, see [Using Extensible Attributes](#).
4. Save the configuration.

## Modifying SNMPv3 Users

1. From the **Administration** tab, select the **SNMPv3 Users** tab -> *snmpv3user*, and then click the Edit icon.
2. The *SNMPv3 User* editor provides the following tabs from which you can edit data:
  - **General:** Modify the data as described in *Configuring SNMPv3 Users* in this topic.
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with the SNMPv3 user account. You can also modify the values of extensible attributes. For information, see [Using Extensible Attributes](#).
3. Save the configuration.

## Deleting SNMPv3 Users

When you delete an SNMPv3 user that is configured to send queries or receive traps, a warning message states that the SNMPv3 is associated with the corresponding function. You can then decide whether you want to delete the user or not. To delete an SNMPv3 user:

1. From the **Administration** tab, select the **SNMPv3 Users** tab -> *snmpv3user*, and then click the Delete icon.
2. In the *Delete confirmation* dialog box, click **Yes**. You cannot schedule the deletion of an SNMPv3 user.

## Accepting Queries

You can allow specific management systems to send SNMP queries to a NIOS appliance. For SNMPv1 and SNMPv2, you must specify a community string. The appliance accepts queries only from management systems that provide the correct community string. You can also specify SNMPv3 users to send queries.

To configure an appliance to accept SNMP queries:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.  
Member: From the **Grid** tab, select the **Grid Manager** -> **Members** tab -> *member*, and then click the **Edit** icon.
2. In the *Grid Properties* or *Grid Member Properties* editor, select the **SNMP** tab. To override Grid settings, click **Override** in the *Grid Member Properties* editor.
3. Complete the following in the **SNMP** section.



- **Enable SNMPv1/SNMPv2 Queries:** Select this to accept SNMPv1 and SNMPv2 queries from management systems.
  - **Community String:** Enter a text string that the management system must send together with its queries to the appliance. A community string is similar to a password in that the appliance accepts queries only from management systems that send the correct community string. Note that this community string must match exactly what you enter in the management system.
- **Engine ID:** Displays the engine ID of the appliance that manages the SNMP agent. The management system needs this ID to send traps to the appliance. If the appliance is an HA pair, this field displays the engine IDs for both the active and passive nodes.
- **Enable SNMPv3 Queries:** Select this to enable queries from SNMPv3 management systems. Click the Add icon to add SNMPv3 users that you have configured on the appliance. In the *SNMPv3 User Selector* dialog box, click the SNMPv3 user you want to add. The appliance displays the selected SNMPv3 users in the table. You can add comments in the table. You can also select an SNMPv3 user and click the Delete icon to remove it from the table. Note that a disabled SNMPv3 user cannot send queries to the appliance.

4. Save the configuration.

## Adding Trap Receivers

You can enable the NIOS appliance to send traps to specific management systems using either SNMPv1/SNMPv2 or SNMPv3, or all versions of SNMP. You can then add management systems that are allowed to receive traps from the appliance. Note that you cannot enable both SNMPv1/SNMPv2 and SNMPv3 on the same trap receiver. The appliance sends traps when certain events occur. You can enable SNMP traps and add trap receivers to the Grid. You can also override the Grid settings at the member level.

To enable the appliance to send traps and to add trap receivers, do the following:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.  
Member: From the **Grid** tab, select the **Grid Manager** -> **Members** tab -> *member*, and then click the **Edit** icon.
2. In the *Grid Properties* or *Grid Member Properties* editor, select the **SNMP** tab. To override Grid settings, click **Override** in the *Grid Member Properties* editor.
3. Complete the following in the **SNMP** tab:
  - **Enable SNMPv1/SNMPv2 Traps:** Select this to enable the appliance to send traps to specified management systems.
    - **Community String:** Enter a text string that the NIOS appliance sends to the management system together with its traps. Note that this community string must match exactly what you enter in the management system.
  - **Enable SNMPv3 Traps:** Select this to enable the appliance to send traps to specified SNMPv3 users.
4. Click the Add icon and select one of the following from the drop-down menu to add an SNMP trap receiver:
  - **SNMPv1/SNMPv2:** Select this to add an SNMPv1 or SNMPv2 management system as a trap receiver. Grid Manager adds a row to the table. In the **Address** field, enter the IP address of the SNMP management system to which you want the SNMP agent on the appliance to send traps. You can enter more than one trap receiver. To remove a trap receiver from the list, select the address, and then click the Delete icon.
  - **SNMPv3:** Select this to add an SNMPv3 management system as a trap receiver. Grid Manager displays the *SNMPv3 User Selector* dialog box. Click the name of the SNMPv3 user in the dialog box. Grid Manager adds the user to the table. In the **Address** field, enter the IP address of the SNMP management system to which you want the SNMP agent on the appliance to send traps. You can add more than one trap receiver. To remove a trap receiver from the list, select the address, and then click the Delete icon. Trap receiver IP addresses may be in IPv4 or IPv6 format.
  - In the Trap Receiver table, Grid Manager displays the following information about the trap receivers:
    - **Address:** The IPv4 or IPv6 address of the trap receiver. Note that when an SNMPv3 user is disabled, SNMPv1/SNMPv2 traps are disabled. You can modify the IP address of the trap receiver even when the following are disabled: SNMPv3 users, SNMPv1/SNMPv2 traps, and SNMPv3 traps.
    - **SNMPv3 User:** The user name of the SNMPv3 trap receiver. This is for SNMPv3 only.
    - **Comment:** Information you entered about the management system.
5. Save the configuration.

## Setting SNMP System Information

You can enter values for certain managed objects in MIB-II, the standard MIB defined in *RFC 1213*. Management systems that are allowed to send queries to the appliance can query these values. You can enter these values for the Grid and specific members. You can also override the Grid values at a member level.

To enter system information:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.  
Member: From the **Grid** tab, select the **Grid Manager** -> **Members** tab -> *member*, and then click the **Edit** icon.
2. In the *Grid Properties* or *Grid Member Properties* editor, select the **SNMP** tab. To override Grid settings, click **Override** in the *Grid Member Properties* editor.
3. Complete the following in the **SNMP** tab. For an HA member, click **Override Node 2 settings** to enter information for node 2 of the HA pair.
  - **sysContact**: Enter the name of the contact person for the appliance.
  - **sysLocation**: Enter the physical location of the appliance.
  - **sysName**: Enter the fully qualified domain name of the appliance.
  - **sysDescr**: Enter useful information about the appliance, such as the software version it is running.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Defining Thresholds for Traps

Threshold events for appliance performance are configurable. For each event, you can set a value that triggers the appliance to send a trap and another value at which the appliance sends a CLEAR trap. The appliance sends a CLEAR trap the first time the event value reaches the reset value after it reached the trigger value.

To define the threshold values:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.  
Member: From the **Grid** tab, select the **Grid Manager** -> **Members** tab -> *member*, and then click the **Edit** icon.
2. In the *Grid Properties* or *Grid Member Properties* editor, click **Toggle Advanced Mode**, and then select the **SNMP Threshold** tab. To override Grid settings, click **Override** in the *Grid Member Properties* editor.
3. Complete the following in the **SNMP Threshold** tab. Each of the following event types have default Trigger and Reset values. You can change the values for any of them. You can set SNMP thresholds above or below which the appliance sends SNMP traps and email notifications, if configured to do so. When any allocated usage exceeds the Trigger value, the appliance sends an SNMP trap and email notification to the designated destination, and the status icon for that usage turns red. When usage drops to the Reset value, the status color goes back to normal and turns green.
  - **CPU Usage**: The percentage of the CPU that is currently in use. The default Trigger value is 81%, and the default Reset value is 70%. Note that these default values are set to disable the CPU usage trap. You can enable this trap and configure the trigger and reset values using the CLI command `set thresholdtrap`. Sometimes you may see high CPU usage on Network Insight discovery members. In most cases, this high CPU usage is the result of CPU-intensive tasks run by the MySQL database during the consolidation of customer data. Network Insight is not a real-time system, so high CPU usage itself does not lead to loss of functionality. Therefore this alert should not be treated as critical if no other negative symptoms are observed on discovery members.
  - **Database Objects**: The percentage of database capacity that is currently in use. The default Trigger value is 80%, and the default Reset value is 70%.
  - **Disk**: The percentage of the primary hard disk that is currently in use. The default Trigger value is 85%, and the default Reset value is 70%.
  - **File Distribution Usage**: The percentage of the file distribution storage capacity that is currently in use on the selected member. The default Trigger value is 90%, and the default Reset value is 70%.
  - **IPAM Utilization**: For a network, this is the percentage based on the IP addresses in use divided by the total addresses in the network and for a network container that contains subnets, this is the percentage of the total address space defined within the container regardless of whether any of the IP addresses in the subnets are in use. The default Trigger value is 95% and the default Reset value is 85%. The status icon turns red when utilization crosses the configured trigger value. When utilization is below the trigger value, the status color turns blue.
  - **Memory**: The percentage of the system memory that is currently in use. The default Trigger value is 90%, and the default Reset value is 80%.

- **Network Capacity:** When the Grid is part of a Master Grid, this is the percentage of the Master Grid's network capacity that is used by the Grid's networks. The default Trigger value is 85% and default Reset value is 75%.
- **Recursive Clients:** The percentage of the limit of concurrent recursive queries. The default Trigger value is 80%, and the default Reset value is 30%. You must also enable the recursive client limit in order for the appliance to send recursive client traps. For information about how to set this limit and restricting recursive client queries, see [Enabling Recursive Queries](#). When you configure the Trigger and Reset values, ensure that you do not set them too low or too close together. If the Trigger and Reset values are too close together, the appliance may send excessive traps and email notifications because both trigger and reset traps are sent based on the calculated value of simultaneous recursive client queries. For example, when you set the recursive client limit at 50, Trigger value at 71%, and Reset value at 70%, the value for simultaneous recursive client queries is calculated at  $50 \times .71 = 35$  (integer math truncation) and  $50 \times .70 = 35$ . This could result in the appliance sending trigger and reset traps for the same value of simultaneous recursive client queries.
- **Root File System:** The percentage of the root file system ("/") that is currently in use. The default Trigger value is 85%, and the default Reset value is 70%.
- **Swap Usage:** The percentage of the swap area that is currently in use. The factory default trigger value is 50% and the factory default reset value is 30%. Grid Manager displays zero for both the trigger and reset values indicating the optimized usage of platform specific default values. For information about available memory on each appliance model, see the table below.
- **Reporting:** The number of reports created on the system that can trigger an SNMP trap. The default Trigger value is 85, and the default Reset value is 70. Note that the maximum number of reports supported per Grid is 300. This field is displayed only when you have configured a reporting server.
- **Reporting Volume:** The percentage of data transmissions to the reporting server. The default Trigger value is 80%, and the default Reset value is 71%. This field is displayed only when you have configured a reporting server.
- **Threat Protection Dropped Traffic:** The percentage of packets dropped based on the threat protection rule configuration. The default Trigger value is 90%, and the default Reset value is 70%. This field is displayed only when Threat Protection licenses are installed on the appliance. When the percentage of Threat Protection dropped traffic exceeds the Trigger value or drops below the Reset value, the appliance sends an SNMP trap and an email notification — if configured to do so. For information about setting SNMP traps and email notifications, see [Setting SNMP and Email Notifications](#) in the following section.
- **Threat Protection Total Traffic:** The percentage of total traffic received (dropped and passed packets) on the external interfaces. The default Trigger value is 90%, and the default Reset value is 70%. This field is displayed only when Threat Protection licenses are installed on the appliance. When the percentage of total Threat Protection traffic exceeds the Trigger value or drops below the Reset value, the appliance sends an SNMP trap and an email notification — if configured to do so.

If you have installed Threat Protection licenses on the appliance and are using the Infoblox feature, Grid Manager displays the following for **Trigger events per second** and **Reset events per second**:

- **Alert Rate:** The number of SNMP traps sent per second when the appliance sends alerts while passing packets based on threat protection rule configuration. The default trigger value is based on the events per second configured for the rule.
- **Drop Rate:** The number of SNMP traps sent per second when the appliance drops packets based on the threat protection rule configuration. The default trigger value resets based on the number of events that are reset.

If you have installed an RPZ license on the NIOS appliance, you can configure the thresholds to monitor the RPZ hit rate in the **Response Policy Zones Hit Rate Configuration** section. For information, see [Configuring Thresholds for RPZ Hit Rate](#).

4. Save the configuration and click **Restart** if it appears at the top of the screen.

For information on available Infoblox appliance models and their available memory see the NIOS 8.6 Release Notes on the Infoblox Support Site. Note that TE appliances are IB appliances in Grid Manager.

## Setting SNMP and Email Notifications

You can specify the event types that trigger trap and email notifications. To set SNMP trap and email notifications:

1. Grid: From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.  
Member: From the **Grid** tab, select the **Grid Manager** -> **Members** tab -> *member*, and then click the **Edit** icon.
2. In the *Grid Properties* or *Grid Member Properties* editor, click **Toggle Advanced Mode**, and then select the **Notification** tab. To override Grid settings, click **Override** in the *Grid Member Properties* editor.
3. Complete the following:
  - **Enable All SNMP Notifications:** Select this checkbox if you want the appliance to send SNMP notifications (traps) for all events to the configured trap receivers. This is selected by default. To send SNMP notifications for specific events to the configured trap receiver, select the checkbox for respective event type.  
For information on configuring trap receivers, see *Adding Trap Receivers* in this topic.
  - **Enable All Email Notifications:** Select this checkbox if you want the appliance to send email notifications (traps) for all events to the configured email recipients. This is deselected by default. To send email notifications for specific events to the configured email recipients, select the checkbox for each respective event type. For information on enabling email notifications and specifying recipients, see [Notifying Administrators](#).
  - Alternatively, you can select specific event types from the table, and specify whether you want the appliance to send SNMP Notifications and Email notifications for each type of event.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

### Selecting SNMP and Email Notification Types

Instead of enabling and receiving SNMP and email notifications for all hardware and software events, you can select specific alert types when a specific hardware or software encounters issues. To enable SNMP or email notifications for specific event types, select the corresponding checkboxes in the Notification tab of the Grid Properties or Member Properties editor. The following table lists the event types you can select:

#### Event Types

Event Type	Description	Sample SNMP Trap
Automated Traffic Capture	Sends notifications each time traffic capture is enabled or disabled or a support bundle is downloaded. For more information, see <a href="#">Enabling Automated Traffic Capture</a> .	01:09:57.938095 IP (tos 0x0, ttl 60, id 37373, offset 0, flags [DF], proto UDP (17), length 356) 10.34.172.4.54004 > 10.120.20.61.162: [udp sum ok] { SNMPv2c { V2Trap(309) R=1114072091 .1.3.6.1.2.1.1.3.0=5985 .1.3.6.1.6.3.1.1.4.1.0=.1.3.6.1.4.1.7779.3.1.1.1.1.7 .1.3.6.1.4.1.7779.3.1.1.2.1.0="infoblox.localdomain" .1.3.6.1.4.1.7779.3.1.1.2.2.0=2 .1.3.6.1.4.1.7779.3.1.1.2.5.0="Automated Traffic Capture" .1.3.6.1.4.1.7779.3.1.1.2.4.0=4 .1.3.6.1.4.1.7779.3.1.1.2.11.0="Automated traffic capture triggered by hitting Queries per second threshold: threshold=100, current=0" } }
BGP	Sends notifications when the BGP software has failed. For more information about the table <i>ibProbableCause Values</i> (OID 3.1.1.1.2.4.0), see <a href="#">SNMP MIB Hierarchy</a> .	2012-11-22 04:49:06 eng-lab-883.inca.infoblox.com [UDP: [10.35.3.115]:38185->[10.120.20.160]]: DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (45366) 0:07:33.66 SNMPv2-MIB::snmpTrapOID.0 = OID: IB-TRAP-MIB::ibTrapOneModule.2 IB-TRAP-MIB::ibName.0 = STRING: "10.35.3.115"  IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: critical(5) IB-TRAP-MIB::ibSubsystemName.0 = STRING: bgp  IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibRevokedLicense(53) IB-TRAP-MIB::ibTrapDesc.0 = STRING: An BGP routing daemon failure has occurred.

Event Type	Description	Sample SNMP Trap
Backup	Sends notifications about the status of backup operation. For more information about the Processing and Software Failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2011-09-22 09:14:17  ib-10-34-41-4.infoblox.com [UDP: [10.34.41.4]:41243-&gt;[10.34.41.4]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (48621) 0:08:06.21  SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibProcessingFailureTrap  IB-TRAP-MIB::ibNodeName.0 = STRING: "10.34.41.4"    IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)  IB-TRAP-MIB::ibSubsystemName.0 = STRING:  scheduled_tftp_backups  IB-TRAP-MIB::ibProbableCause.0 = INTEGER:  ibBackupSoftwareFailure(29)  IB-TRAP-MIB::ibTrapDesc.0 = STRING: Backup failed.</p>
BloxTools	Sends notifications about the status of bloxTools. For more information about the Object State Change traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2011-09-13 20:38:46  10.34.42.4 [UDP: [10.34.42.4]:38187-&gt;[10.34.42.2]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (742156) 2:03:41.56  SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibStateChangeEvent  IB-TRAP-MIB::ibNodeName.0 = STRING: "10.34.42.4"  IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)  IB-TRAP-MIB::ibName.0 = STRING: bloxTools  IB-TRAP-MIB::ibPreviousState.0 = INTEGER: bloxtools-service-failed(41)  IB-TRAP-MIB::ibCurrentState.0 = INTEGER: bloxtools-service-working(39)  IB-TRAP-MIB::ibTrapDesc.0 = STRING: BloxTools Service is working.</p>
CPU	Sends notifications about the status of CPU usage. For more information about the Threshold Crossing Event traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2012-04-12 01:54:59  eng-lab-631.inca.infoblox.com [UDP: [10.35.2.119]:42546-&gt;[10.120.20.160]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (786308) 2:11:03.08  SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibTrapOneModule.3  IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.2.119"    IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: major(4)  IB-TRAP-MIB::ibName.0 = STRING: cpu_usage  IB-TRAP-MIB::ibCurThresholdValue.0 = INTEGER: 51    IB-TRAP-MIB::ibThresholdHigh.0 = INTEGER: 5  IB-TRAP-MIB::ibThresholdLow.0 = INTEGER: 3    IB-TRAP-MIB::ibTrapDesc.0 = STRING: CPU usage above threshold value.</p>

Event Type	Description	Sample SNMP Trap
CaptivePortal	Sends notifications about the Captive Portal service. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-08 02:58:23</p> <p>10.35.107.4 [UDP: [10.35.107.4]:45111-&gt;[10.120.20.21]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (22280) 0:03:42.80</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.107.4"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: captive_portal</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: captive-portal-service-inactive(51)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: captive-portal-service-working(49)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Captive Portal Service is working.</p>
Cisco ISE	Sends notifications about the status of the Cisco ISE service. For more information about Object State Change Traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-07 22:26:19</p> <p>10.40.240.111 [UDP: [10.40.240.111]:47355-&gt;[10.120.20.21]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (16384) 0:02:43.84</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.40.240.111"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: cisco_ise_server</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: 124</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: 123</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: CiscoISE server 10.36.141.15 is OK.</p>
Clear	Sends notifications when the SNMP trap is cleared. When you select the checkbox, the CLEAR trap is sent for the following software failures: LDAP servers, OCSP responders, LCD, Serial Console, OSPF, OSPF6, BGP, HSM, Controld, SSH, HTTP, Cluster, Login, and Duplicate IP. For file distribution, the trap is sent when the service is restored. If you deselect the checkbox, the CLEAR trap is not sent when any of the mentioned software fails. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2015-11-12 22:46:21</p> <p>10.35.124.1 [UDP: [10.35.124.1]:48446-&gt;[10.120.20.21]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (12793) 0:02:07.93</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.124.1"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: Login</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibClear(0)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: SNMP Trap is cleared. GUI Login</p>

Event Type	Description	Sample SNMP Trap
Cloud API	Sends notifications about whether the Cloud API service is functioning or not. For information about Cloud Network Automation, see Introduction to <a href="#">Cloud Network Automation</a> .	<p>2014-11-13 00:52:37</p> <p>10.35.114.10 [UDP: [10.35.114.10]:57772-&gt;[10.120.20.21]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (480221) 1:20:02.21</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.114.10"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: cloud_api</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: 105</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: 103</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Cloud API service is working.</p>



Event Type	Description	Sample SNMP Trap
Cloud DNS Sync	Sends notification about the status of the Cloud DNS Sync service.	DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (45114) 0:07:31.14 SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.7779.3.1.1.1.4 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.1.0 = STRING: "172.29.2.154" SNMPv2-SMI::enterprises.7779.3.1.1.1.2.2.0 = INTEGER: 2 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.3.0 = STRING: "cloud_dns_sync" SNMPv2-SMI::enterprises.7779.3.1.1.1.2.9.0 = INTEGER: 166 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.10.0 = INTEGER: 169 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.11.0 = STRING: "Cloud DNS Sync Service is not running."  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (50781) 0:08:27.81 SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.7779.3.1.1.1.4 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.1.0 = STRING: "172.29.2.154" SNMPv2-SMI::enterprises.7779.3.1.1.1.2.2.0 = INTEGER: 2 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.3.0 = STRING: "cloud_dns_sync" SNMPv2-SMI::enterprises.7779.3.1.1.1.2.9.0 = INTEGER: 169 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.10.0 = INTEGER: 167 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.11.0 = STRING: "Cloud DNS Sync Service is unhealthy."  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (52645) 0:08:46.45 SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.7779.3.1.1.1.4 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.1.0 = STRING: "172.29.2.154" SNMPv2-SMI::enterprises.7779.3.1.1.1.2.2.0 = INTEGER: 2 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.3.0 = STRING: "cloud_dns_sync" SNMPv2-SMI::enterprises.7779.3.1.1.1.2.9.0 = INTEGER: 167 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.10.0 = INTEGER: 168 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.11.0 = STRING: "Cloud DNS Sync Service is initializing."  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (55591) 0:09:15.91 SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.7779.3.1.1.1.4 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.1.0 = STRING: "172.29.2.154" SNMPv2-SMI::enterprises.7779.3.1.1.1.2.2.0 = INTEGER: 2 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.3.0 = STRING: "cloud_dns_sync" SNMPv2-SMI::enterprises.7779.3.1.1.1.2.9.0 = INTEGER: 168 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.10.0 = INTEGER: 166 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.11.0 = STRING: "Cloud DNS Sync Service is healthy."

Event Type	Description	Sample SNMP Trap
Cluster	Sends notifications about the status of NIOS clusterd process. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2011-12-10 09:43:23</p> <p>infoblox.localdomain [UDP: [10.35.2.70]:44193-&gt;[10.35.2.70]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (47140) 0:07:51.40</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "0"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: critical(5)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: Clusterd_Monitor</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibFDSsoftwareFailure(32)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: WATCHDOG: A grid daemon failure has occurred on 10.35.2.70</p>
Controld	Sends notifications about the NIOS controld process. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2012-08-17 05:29:30</p> <p>&lt;UNKNOWN&gt; [UDP: [10.32.2.80]:43475-&gt;[10.32.2.80]:162]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (24722) 0:04:07.22</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "192.168.1.2"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: critical(5)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: controld</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibControldSoftwareFailure(11)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: A controld failure has occurred.</p>
DHCP	Sends notifications about the status of DHCP service. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-04-18 02:42:36</p> <p>10.35.139.15 [UDP: [10.35.139.15]:35531-&gt;[10.120.21.204]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (8809) 0:01:28.09</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.139.15"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: dhcpd</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: dhcp-service-inactive(48)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: dhcp-service-working(45)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: DHCP Service is working.</p>

Event Type	Description	Sample SNMP Trap
DNS	Sends notifications about the status of DNS service. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-08 01:10:53</p> <p>10.35.3.154 [UDP: [10.35.3.154]:59876-&gt;[10.120.20.12]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (324160) 0:54:01.60</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.3.154"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING:DNS</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: dns-service-working(32)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: dns-service-inactive(34)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: DNS Service is inactive.</p>
DNS Attack	Sends notifications about the status of the DNS attacks. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-07 23:55:56</p> <p>10.35.3.201 [UDP: [10.35.3.201]:33199-&gt;[10.120.21.204]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (25100) 0:04:11.00</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.3.201"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: DNSAttack</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: dns-attack-active(115)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: dns-attack-inactive(116)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: DNS attack conditions have ended.</p>
DNS Forwarding to BloxOne Threat Defense Cloud	Sends notifications about whether DNS forwarding to BloxOne Threat Defense Cloud is functioning or not. For information about DNS forwarding to BloxOne Threat Defense Cloud, see <a href="#">Using Forwarders</a> .	<p>2017-11-03 14:06:19 192.168.1.3 [UDP: [192.168.1.3]:33577-&gt;[192.168.1.2]:162]: Trap, DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (180166) 0:30:01.66, SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>SNMPv2-SMI::enterprises.7779.3.1.1.1.4, SNMPv2-SMI::enterprises.7779.3.1.1.2.1.0=STRING: "192.168.1.3", SNMPv2-SMI::enterprises.7779.3.1.1.2.2.0 =INTEGER: 2, SNMPv2-SMI::enterprises.7779.3.1.1.2.3.0 =STRING: "DNS", SNMPv2-SMI::enterprises.7779.3.1.1.2.9.0 =INTEGER: 32, SNMPv2-SMI::enterprises.7779.3.1.1.2.10.0= INTEGER: 133, SNMPv2-SMI::enterprises.7779.3.1.1.2.11.0= STRING: "The appliance is still serving DNS even though forwarding DNS queries to BloxOne Threat Defense Cloud is not functioning properly.</p>

Event Type	Description	Sample SNMP Trap
DNS Integrity Check	Sends notifications about whether DNS integrity check is functioning or not. For information about DNS integrity check, see <a href="#">Configuring DNS Integrity Check for Authoritative Zones</a> .	<p>2014-06-03 05:35:45</p> <p>10.34.82.121 [UDP: [10.34.82.121]:42577-&gt;[10.120.20.232]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance =</p> <p>Timeticks: (427200) 1:11:12.00</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.34.154.3"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: DNS Integrity Check</p> <p>IB-TRAP-MIB::ibProbableCause.0 =INTEGER: ibDNSIntegrityCheckNameserversFailed(102)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: DNS Integrity Check is unable to connect to any name servers required to perform the check. There are list of nameservers failure: ([10.35.0.56'])</p>
DNS Integrity Check Connection	Sends notifications about whether DNS integrity check connection is functioning or not. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>016-01-12 05:10:27</p> <p>10.35.129.15 [UDP: [10.35.129.15]:35201-&gt;[10.120.20.12]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (213856) 0:35:38.56</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.129.15"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0= INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: DNS Integrity Check</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: dns-integrity-check-severity-indetermined(97)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: dns-integrity-check-severity-informational( 99)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: DNS Integrity Check severity has changed. For DNS zone 'default-com.info', severity has changed from 'NONE' to 'INFORMATIONAL'. Reason: While checking server ib-10-35-133-2.infoblox.com.(ipv4=10.35.133.2) INFORMATIONAL discrepancy found because our servers have record(s): info.com. I N NS ib-10-35-12</p>

Event Type	Description	Sample SNMP Trap
Database	Sends notifications about the database status. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2013-04-05 02:27:01</p> <p>10.35.116.2 [UDP: [10.35.116.2]:45332-&gt;[10.120.20.160]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (120021) 0:20:00.21</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibThresholdCrossingEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.116.2"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: db_usage</p> <p>IB-TRAP-MIB::ibCurThresholdValue.0 = INTEGER: 1</p> <p>IB-TRAP-MIB::ibThresholdHigh.0 = INTEGER: 85</p> <p>IB-TRAP-MIB::ibThresholdLow.0 = INTEGER: 0</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Database capacity used is OK.</p>
Disconnected Grid	Sends notifications about whether a Grid has been disconnected from the Master Grid. For more information about object state change traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2011-12-27 23:38:44</p> <p>eng-lab-089.inca.infoblox.com [UDP: [10.35.0.89]:53010-&gt;[10.120.20.160]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (1049) 0:00:10.49</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.0.89"</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: ID_Grid</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: grid-disconnected(5)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: grid-connected(4)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: The grid member is connected to the grid master.</p>
Discovery	Sends notifications about the discovery status. For information about the Discovery feature, see <a href="#">Infoblox Network Insight</a> .	<p>2013-10-22 01:35:53</p> <p>eng-lab-302.inca.infoblox.com [UDP: [10.35.1.46]:57126-&gt;[10.120.20.102]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (122717) 0:20:27.17</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.1.46"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: discovery_collector</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: 89</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: 86</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Discovery Collector Service is working.</p>

Event Type	Description	Sample SNMP Trap
Discovery Conflict	Sends notifications about conflicts between the DHCP address and the existing IP address. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2014-06-20 00:48:06</p> <p>10.34.125.2 [UDP: [10.34.125.2]:36159-&gt;[10.120.20.232]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (2241) 0:00:22.41</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.34.125.2"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: major(4)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: DHCP</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibDHCPHostConflict(66)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: DHCP address conflicts with an existing host address. [<i>IP address</i>]</p>
Discovery Unmanaged	Sends notifications related to the discovery of unmanaged devices and networks. You can configure the maximum number of unmanaged objects the appliance discovers and how often it notifies about these events. For more information about how to configure these parameters, see <a href="#">Defining Seed Routers for Probe Members</a> .	<p>2015-02-09 22:53:57 10.35.103.18 [UDP: [10.35.103.18]:45156-&gt;[10.120.20.46]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (46090) 0:07:40.90</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibOperationTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.103.18"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: major(4)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: Discovery</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: 4012</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: New unmanaged devices/networks were found during network discovery process. New unmanaged devices in '10.40.16.0/20' network in 'default' network view.</p>
Disk	Sends notifications about the status of the primary disk. For more information about threshold crossing traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2012-11-22 03:34:32</p> <p>eng-lab-883.inca.infoblox.com [UDP: [10.35.3.115]:37542-&gt;[10.120.20.160]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (87141) 0:14:31.41</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibTrapOneModule.3</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.3.115"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: disk_usage</p> <p>IB-TRAP-MIB::ibCurThresholdValue.0 = INTEGER: 15</p> <p>IB-TRAP-MIB::ibThresholdHigh.0 = INTEGER: 10</p> <p>IB-TRAP-MIB::ibThresholdLow.0 = INTEGER: 5</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: System primary hard disk usage is over threshold value.</p>

Event Type	Description	Sample SNMP Trap
DuplicateIP	Sends notifications when there are duplicate IP addresses. For more information about <code>ibProbableCause</code> Values, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-04-18 22:47:38</p> <p>10.35.139.15 [UDP: [10.35.139.15]:35531-&gt;[10.120.21.204]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (7239201) 20:06:32.01</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.139.15"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: major(4)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: Equipment</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibDuplicateIPAddressFailure(52)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: A Duplicate IP Address has been detected.</p>
ENAT	Sends notifications about the Ethernet port status. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2012-08-27 03:05:25</p> <p>10.36.3.132 [UDP: [10.36.3.132]:47962-&gt;[10.120.20.232]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (16624) 0:02:46.24</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.36.3.132"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: major(4)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: LAN</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: enet-link-up(6)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: enet-link-down(7)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: LAN1 port link is down. Please check the connection.</p>
File Distribution Usage	Sends notifications about the HTTP file distribution process. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-08 01:48:57</p> <p>10.40.240.113 [UDP: [10.40.240.113]:41443-&gt;[10.120.20.21]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (183757) 0:30:37.57</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.40.240.113"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: critical(5)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: fd_usage</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibFDSoftwareFailure(42)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: File Distribution services storage usage is OK.</p>



Event Type	Description	Sample SNMP Trap
FTP	Sends notifications about the status of FTP service. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-07 23:27:22</p> <p>10.40.240.113 [UDP: [10.40.240.113]:36063-&gt;[10.120.20.21]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (380473) 1:03:24.73</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.40.240.113"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: critical(5)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING:</p> <p>ftp IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibFTPDSoftwareFailure(44)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: An FTPD daemon failure has occurred.</p>
Fan	Sends notifications about the status of the system fan. For more information about IEquipment Failure Traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2012-02-23 23:34:50</p> <p>10.32.1.222 [UDP: [10.32.1.222]:42742-&gt;[10.35.109.24]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (89508) 0:14:55.08</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibEquipmentFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.32.1.222"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: fan</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibFan1Failure(37)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Fan 1 failure has occurred.</p>
HA	Sends notifications about the status of the HA port link. For more information about the ibPreviousState (OID 3.1.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2014-05-23 00:49:32</p> <p>eng-lab-589.inca.infoblox.com [UDP: [10.35.2.77]:46426-&gt;[10.36.0.200]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (3912) 0:00:39.12</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.20.70"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: HA_Replication</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: ha-replication-offline(14)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: ha-replication-online(13)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: HA replication is online.</p>

Event Type	Description	Sample SNMP Trap
HSM	Sends notifications about the status of the HSM operation. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-12 20:52:12</p> <p>10.39.13.77 [UDP: [10.39.13.77]:44962-&gt;[10.120.20.21]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (131909) 0:21:59.09</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.39.13.77"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: DNS</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: dns-service-working(32)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: 55</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: DNS Service restarted not using HSM</p>
HTTP	Sends notifications about the status of the HTTP service. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-07 23:16:40</p> <p>10.40.240.113 [UDP: [10.40.240.113]:36063-&gt;[10.120.20.21]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (316247) 0:52:42.47</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.40.240.113"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: http_file_dist</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: http-file-dist-service-inactive(40)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: http-file-dist-service-working(38)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: HTTP File Dist Service is working</p>
IFMAP	Sends notifications about the status of the IF-MAP service. For more information about <code>ibProbableCause</code> Values, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-04-20 23:40:00</p> <p>10.34.41.60 [UDP: [10.34.41.60]:33231-&gt;[10.120.21.204]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (45256) 0:07:32.56</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.34.41.60"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: critical(5)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: ifmapd</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibIFMAPSoftwareFailure(50)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: An IF-MAP server failure has occurred.</p>

Event Type	Description	Sample SNMP Trap
IPMI Device	Sends notifications about the status of the IPMI device. For more information about <code>ibProbableCause</code> Values, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2015-03-24 04:10:21 eng-lab-598.inca.infoblox.com [UDP: [10.35.2.86]:56092-&gt;[10.120.20.21]]: DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (103490) 0:17:14.90 SNMPv2-MIB::snmpTrapOID.0 = OID: IB-TRAP-MIB::ibOperationTrap IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.2.86"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2) IB-TRAP-MIB::ibSubsystemName.0 = STRING: IPMI Status Check IB-TRAP-MIB::ibProbableCause.0 = INTEGER: 2105</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: IPMI is used by some hardware monitors to test hardware health. The IPMI Device is now available; subsequent hardware monitor failures are likely to be genuine.</p>
IPAM Utilization	Sends notifications about the percentage of IPv4 addresses that are used in a network. For a network container that contains subnets, this indicates the percentage of the total address space defined within the container regardless of whether any of the IP addresses are used in the subnetwork. For more information about threshold crossing traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2014-07-06 19:33:01 eng-lab-514.inca.infoblox.com [UDP: [10.35.2.2]:33413-&gt;[10.120.20.21]]: DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (1038846) 2:53:08.46 SNMPv2-MIB::snmpTrapOID.0 = OID: IB-TRAP-MIB::ibThresholdCrossingEvent IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.2.2"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: Threshold IB-TRAP-MIB::ibCurThresholdValue.0 = INTEGER: 0</p> <p>IB-TRAP-MIB::ibThresholdHigh.0 = INTEGER: 5 IB-TRAP-MIB::ibThresholdLow.0 = INTEGER: 3</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Network IPAM Utilization capacity usage is OK. Network: 20.0.0.0/29/netview</p>
Load Balancer Device	Sends notifications about whether the LB device is in sync or not. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2013-07-08 02:36:39 10.35.113.2 [UDP: [10.35.113.2]:45568-&gt;[10.120.20.21]]: DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (996274) 2:46:02.74</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID: IB-TRAP-MIB::ibStateChangeEvent IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.113.2"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2) IB-TRAP-MIB::ibObjectName.0 = STRING: device IB-TRAP-MIB::ibPreviousState.0 = INTEGER: lb-device-down(74) IB-TRAP-MIB::ibCurrentState.0 = INTEGER: lb-device-up(73)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Load Balancer device 10.36.128.1 sync is OK.</p>

Event Type	Description	Sample SNMP Trap
LCD	Sends notifications about the status of the LCD process. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2011-12-01 06:31:28  ib-10-35-3-125.infoblox.com [UDP: [10.35.3.125]:56609-&gt;[10.35.3.125]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (187334) 0:31:13.34  SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibProcessingFailureTrap  IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.3.125"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: major(4)  IB-TRAP-MIB::ibSubsystemName.0 = STRING: lcd</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER:  ibLCDSoftwareFailure(18)  IB-TRAP-MIB::ibTrapDesc.0 = STRING: An LCD failure has occurred.</p>
LDAP Servers	Sends notifications about whether LDAP servers are available or not. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2012-12-17 03:33:58  10.35.106.6 [UDP: [10.35.106.6]:35751-&gt;[10.120.20.249]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (4043) 0:00:40.43</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.106.6"  IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)  IB-TRAP-MIB::ibObjectName.0 = STRING: ldap_servers</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: ldap-service-inactive(79)  IB-TRAP-MIB::ibCurrentState.0 = INTEGER: ldap-servers-ok(76)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: All LDAP servers are available.</p>
License	Sends notifications when the license has been revoked. For more information about Revoked License Trap, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2014-05-23 00:49:32  eng-lab-589.inca.infoblox.com [UDP: [10.35.2.77]:46426-&gt;[10.36.0.200]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (3912) 0:00:39.12  SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibTrapOneModule.6.0  IB-TRAP-MIB::ibNodeName.0 = STRING: "10.34.196.165"  IB-TRAP-MIB::ibSubsystemName.0 = STRING: vnios</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Grid license of member IP 10.34.196.165 with hardware ID 564d5cf2115391ab3f6efee9a9d974aa is revoked. 0 Cold Start</p>

Event Type	Description	Sample SNMP Trap
Login	Sends notifications when the login details are incorrect. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2013-04-30 02:56:49</p> <p>10.34.132.2 [UDP: [10.34.132.2]:35453-&gt;[10.120.20.160]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (27604) 0:04:36.04</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.34.152.2"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: Login</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibGUILoginFailure(58)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: A GUI login failure has occurred.</p>
MGM	Sends notifications about the status of the multi-Grid configuration. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2013-04-03 23:48:34</p> <p>eng-lab-482.inca.infoblox.com [UDP: [10.35.1.226]:37893-&gt;[10.120.20.160]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (24697) 0:04:06.97</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibTrapOneModule.4</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.116.2"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: major(4)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: Grid of Grids</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: 23</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: 24</p> <p>IB-TRAP-MIB::ibTrapDesc.0 =STRING: Grid connection offline.</p>
MSServer	Sends notifications about the status of Microsoft Servers for Microsoft management. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2013-07-10 03:29:42</p> <p>10.35.113.2 [UDP: [10.35.113.2]:51679-&gt;[10.120.20.21]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (1066392) 2:57:43.92</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.113.2"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: ms_service</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: ms-service-down(20)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: ms-service-up(19)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Service connection to Microsoft DNS server 10.102.31.67 is OK.</p>

Event Type	Description	Sample SNMP Trap
Memory	Sends notifications about the status of the system memory. For more information about threshold crossing traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2012-04-17 03:09:46</p> <p>10.35.119.4 [UDP: [10.35.119.4]:37664-&gt;[10.120.20.160]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (359835) 0:59:58.35</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibTrapOneModule.3</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.119.4"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: memory</p> <p>IB-TRAP-MIB::ibCurThresholdValue.0 = INTEGER: 46</p> <p>IB-TRAP-MIB::ibThresholdHigh.0 = INTEGER: 50</p> <p>IB-TRAP-MIB::ibThresholdLow.0 = INTEGER: 30</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: System memory usage is OK.</p>
NTP	Sends notifications about the status of the NTP service. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2013-04-02 23:58:19</p> <p>10.35.116.6 [UDP: [10.35.116.6]:37505-&gt;[10.120.20.160]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (17804) 0:02:58.04</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.1.187"</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: ntp_sync</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: ntp-sync-down(16)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: ntp-sync-up(15)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: The NTP service resumed synchronization.</p>
Network	Sends notifications about the status of the LAN port. For more information about threshold crossing traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2013-01-06 23:52:01</p> <p>10.35.3.62 [UDP: [10.35.3.62]:49255-&gt;[10.120.20.160]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (80879) 0:13:28.79</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibTrapOneModule.3</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.3.62"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: MGM</p> <p>IB-TRAP-MIB::ibCurThresholdValue.0 = INTEGER: 8</p> <p>IB-TRAP-MIB::ibThresholdHigh.0 = INTEGER: 5</p> <p>IB-TRAP-MIB::ibThresholdLow.0 = INTEGER: 3</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Network capacity used is over the threshold value.</p>

Event Type	Description	Sample SNMP Trap
OCSP Responders	Sends notifications about the status of OCSP responders. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-14 03:21:14</p> <p>10.34.9.91 [UDP: [10.34.9.91]:34663-&gt;[10.120.21.204]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (4599) 0:00:45.99</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.34.9.91"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2) IB-TRAP-MIB::ibObjectName.0 = STRING: ocspr_responders</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: ocspr_responders-ok(65)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: ocspr-service-inactive(68)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: OCSP service inactive.</p>
OSPF	Sends notifications about the ospfd process. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2012-11-22 04:49:56</p> <p>eng-lab-883.inca.infoblox.com [UDP: [10.35.3.115]:38185-&gt;[10.120.20.160]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (50414) 0:08:24.14</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibTrapOneModule.2</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.3.115"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: critical(5)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: ospf</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibOSPFSoftwareFailure(35)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: An OSPF routing daemon failure has occurred.</p>
OSPF6	Sends notifications about the ospf process for IPv6. For more information about <code>ibProbableCause</code> Values (OID 3.1.1.1.2.4.0), see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-13 02:03:07</p> <p>eng-lab-396.inca.infoblox.com [UDP: [10.35.1.140]:45733-&gt;[10.120.20.21]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (3970) 0:00:39.70</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.1.140"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: ospf6</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibClear(0)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: SNMP Trap is cleared. OSPF6 SW</p>



Event Type	Description	Sample SNMP Trap
PowerSupply	Sends notifications about the status of the power supply. For more information about IEquipment Failure Traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2012-01-20 22:45:01 nextgen.com [UDP: [10.32.111.110]:45323-&gt;[10.32.111.110]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (21044) 0:03:30.44  SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibEquipmentFailureTrap  IB-TRAP-MIB::ibNodeName.0 = STRING: "10.32.111.110"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)  IB-TRAP-MIB::ibObjectName.0 = STRING: power_supply</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER:  ibSystemRestart(61)  IB-TRAP-MIB::ibTrapDesc.0 = STRING: Power supply 2 is OK.</p>
RAID	Sends notifications about the RAID array status. For more information about IEquipment Failure Traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-15 14:30:20  eng-lab-418.inca.infoblox.com [UDP: [10.35.1.162]:57616-&gt;[10.120.21.204]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (8737) 0:01:27.37  SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibEquipmentFailureTrap  IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.1.162"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: major(4)  IB-TRAP-MIB::ibObjectName.0 = STRING: raid  IB-TRAP-MIB::ibProbableCause.0 = INTEGER:  ibRAIDIsDegraded(3002)  IB-TRAP-MIB::ibTrapDesc.0 = STRING: The system's RAID array is in a degraded state.</p>
Recursive Clients	Sends notifications about whether the DNS recursive server is under flood attacks. For more information about threshold crossing traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2015-01-13 02:05:55  eng-lab-078.inca.infoblox.com [UDP: [10.35.0.78]:55233-&gt;[10.120.20.21]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (168817) 0:28:08.17  SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibThresholdCrossingEvent  IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.0.78"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)  IB-TRAP-MIB::ibObjectName.0 = STRING: RecursiveClients  IB-TRAP-MIB::ibCurThresholdValue.0 = INTEGER: 1</p> <p>IB-TRAP-MIB::ibThresholdHigh.0 = INTEGER: 800  IB-TRAP-MIB::ibThresholdLow.0 = INTEGER: 300</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Number of simultaneous DNS recursive lookups is OK.</p>

Event Type	Description	Sample SNMP Trap
RIR SWIP	Sends notifications about the status of the RIR SWIP registration. For more information about <code>ibProbableCause</code> Values (OID 3.1.1.1.2.4.0), see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-11 02:05:17</p> <p>10.34.11.100 [UDP: [10.34.11.100]:52957-&gt;[10.120.21.204]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (43623) 0:07:16.23</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.34.11.100"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: major(4)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: RIR SWIP</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibRIRSWIPRegistrationFailure(89)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: An RIR SWIP registration failure has occurred.CREATE RIR registration request for network "187.0.0.0/24" failed.</p>
Reporting	Sends notifications about the status of the reporting database. For more information about threshold crossing traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2012-01-06 08:59:55 10.35.101.27 [UDP: [10.35.101.27]:59714-&gt;[10.35.117.24]]:DISMA</p> <p>N-EVENT-MIB::sysUpTimeInstance = Timeticks: (289200) 0:48:12.00</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibThresholdCrossingEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.34.46.6"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: reporting_volume</p> <p>IB-TRAP-MIB::ibCurThresholdValue.0 = INTEGER: 85</p> <p>IB-TRAP-MIB::ibThresholdHigh.0 = INTEGER: 80</p> <p>IB-TRAP-MIB::ibThresholdLow.0 = INTEGER: 71</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Reporting volume usage reached the threshold value.</p>
RPZ Hit Rate	Send Notifications about the percentage of RPZ Hit Rate.For more information about threshold crossing traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-04-18 02:52:29</p> <p>10.35.139.15 [UDP: [10.35.139.15]:35531-&gt;[10.120.21.204]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (68173) 0:11:21.73</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibThresholdCrossingEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.139.15"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: major(4)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: RPZHitRate</p> <p>IB-TRAP-MIB::ibCurThresholdValue.0 = INTEGER: 100</p> <p>IB-TRAP-MIB::ibThresholdHigh.0 = INTEGER: 1</p> <p>IB-TRAP-MIB::ibThresholdLow.0 = INTEGER: 0</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: RPZ hit rate is too high: 100%.</p>

Event Type	Description	Sample SNMP Trap
RootFS	Sends notifications about the status of the root file system. For more information about threshold crossing traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2011-11-29 06:36:24  ib-10-35-1-144.infoblox.com [UDP: [10.35.1.144]:59707-&gt;[10.35.1.144]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (197388) 0:32:53.88  SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibThresholdCrossingEvent  IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.1.144"    IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)  IB-TRAP-MIB::ibObjectName.0 = STRING: rootfs_usage  IB-TRAP-MIB::ibCurThresholdValue.0 = INTEGER: 28    IB-TRAP-MIB::ibThresholdHigh.0 = INTEGER: 100    IB-TRAP-MIB::ibThresholdLow.0 = INTEGER: 90    IB-TRAP-MIB::ibTrapDesc.0 = STRING: Root filesystem disk usage is OK.</p>
SNMP	Sends notifications about the status of the SNMP server. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2012-02-17 06:24:11  eng-lab-630.inca.infoblox.com [UDP: [10.35.2.118]:57802-&gt;[10.120.20.174]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (855217) 2:22:32.17  SNMPv2-MIB::snmpTrapOID.0 = OID:  SNMPv2-SMI::enterprises.7779.3.1.1.1.1.2  SNMPv2-SMI::enterprises.7779.3.1.1.1.2.1.0 = STRING: "10.35.2.118"  SNMPv2-SMI::enterprises.7779.3.1.1.1.2.2.0 = INTEGER: 4  SNMPv2-SMI::enterprises.7779.3.1.1.1.2.5.0 = STRING: "snmp"  SNMPv2-SMI::enterprises.7779.3.1.1.1.2.4.0 = INTEGER: 13  SNMPv2-SMI::enterprises.7779.3.1.1.1.2.11.0 = STRING: "SNMP Server failure has occurred."</p>
SSH	Sends notifications about the status of the sshd process. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2011-09-21 21:33:17    10.34.42.6 [UDP: [10.34.42.6]:49776-&gt;[10.34.42.2]]:    DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (6508) 0:01:05.08    SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibProcessingFailureTrap  IB-TRAP-MIB::ibNodeName.0 = STRING: "10.34.42.6"  IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)  IB-TRAP-MIB::ibSubsystemName.0 = STRING: ssh  IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibClear(0)    IB-TRAP-MIB::ibTrapDesc.0 = STRING: SNMP Trap is cleared. SSHD SW</p>

Event Type	Description	Sample SNMP Trap
SerialConsole	Sends notifications when the serial console login has failed or the admin failed to login to the serial console. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2014-12-15 02:47:59</p> <p>UDP/IPV6: 2620:10a:6000:2400::8104]:55038 [UDP/IPV6: [2620:10a:6000:2400::8104]:55038]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (501585) 1:23:35.85</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "2620:10a:6000:2400::8104"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: Login</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibSerialConsoleLoginFailure(59)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: A Serial Console login failure has occurred.</p>
Swap Usage	Sends notifications about whether the swap usage has exceeded the trigger or reset value. For more information about defining thresholds for traps, see above.	<p>2013-11-25 05:35:56 10.35.129.1 [UDP: [10.35.129.1]:49489-&gt;[10.120.20.21]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (24009) 0:04:00.09</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibThresholdCrossingEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.129.1"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: major(4)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: swap_usage</p> <p>IB-TRAP-MIB::ibCurThresholdValue.0 = INTEGER: 100</p> <p>IB-TRAP-MIB::ibThresholdHigh.0 = INTEGER: 5</p> <p>IB-TRAP-MIB::ibThresholdLow.0 = INTEGER: 2</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: System swap space usage exceeds the critical threshold value.</p>
Syslog	Sends notifications when the syslog process stops. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-14 01:29:42</p> <p>10.34.142.105 [UDP: [10.34.142.105]:37566-&gt;[10.120.21.204]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (296285) 0:49:22.85</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibProcessingFailureTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.34.142.105"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: critical(5)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: check_syslog_conf</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: ibSyslogFailure(67)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Syslog daemon is not running.</p>

Event Type	Description	Sample SNMP Trap
System	Sends notifications about the status of the NIOS system. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2014-03-14 05:54:42  eng-lab-636.inca.infoblox.com [UDP: [10.35.2.124]:33232-&gt;[10.36.0.200]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (2654067) 7:22:20.67  SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibProcessingFailureTrap  IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.120.17"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: minor(3)  IB-TRAP-MIB::ibSubsystemName.0 = STRING: ID_Grid</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER:  ibSystemRestart(61)  IB-TRAP-MIB::ibTrapDesc.0 = STRING: The system is being restarted.</p>
TAXII	Sends notifications when you start and stop the TAXII service. For more information about object state change traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-06 20:16:13  eng-lab-242.inca.infoblox.com [UDP: [10.35.0.242]:42213-&gt;[10.120.20.21]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (17553) 0:02:55.53  SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibStateChangeEvent  IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.0.242"  IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)  IB-TRAP-MIB::ibObjectName.0 = STRING: taxii  IB-TRAP-MIB::ibPreviousState.0 = INTEGER: 122  IB-TRAP-MIB::ibCurrentState.0 = INTEGER: 119  IB-TRAP-MIB::ibTrapDesc.0 = STRING: TAXII Service is working.</p>
TFTP	Sends notifications about the status of the TFTP service. For more information about processing and software failure traps, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2011-09-16 06:54:08  eng-lab-443.inca.infoblox.com [UDP: [10.35.1.187]:35794-&gt;[10.120.20.160]]:  DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (50382) 0:08:23.82  SNMPv2-MIB::snmpTrapOID.0 = OID:  IB-TRAP-MIB::ibProcessingFailureTrap  IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.1.187"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: critical(5)  IB-TRAP-MIB::ibSubsystemName.0 = STRING: tftp</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER:  ibTFTPDSoftwareFailure(27)  IB-TRAP-MIB::ibTrapDesc.0 = STRING: A TFTP daemon failure has occurred.</p>

Event Type	Description	Sample SNMP Trap
Threat Analytics	Sends notifications about the status of the Threat Analytics service. For more information about <code>ibProbableCause</code> Values (OID 3.1.1.1.2.4.0), see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-08 00:19:34</p> <p>10.35.3.154 [UDP: [10.35.3.154]:59876-&gt;[10.120.20.12]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (16184) 0:02:41.84</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.3.154"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: Threat Analytics</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: 127</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: 126</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Threat Analytics Service is working.</p>
Threat Analytics DNS Tunneling	Sends notifications about the DNS tunneling detection. For more information about <code>ibProbableCause</code> Values (OID 3.1.1.1.2.4.0), see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-01-08 00:22:07</p> <p>10.35.3.154 [UDP: [10.35.3.154]:59876-&gt;[10.120.20.12]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (31567) 0:05:15.67</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibOperationTrap</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.3.154"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: major(4)</p> <p>IB-TRAP-MIB::ibSubsystemName.0 = STRING: Software</p> <p>IB-TRAP-MIB::ibProbableCause.0 = INTEGER: 4013</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: DNS Tunneling detected: The record '*.any.knjuqljsfyyc2t.com' is successfully added into BlackList RPZ zone with comment '[2016-08-01 08:22:06 UTC] [member: infoblox.localdomain] DNS Tunneling' and policy 'No Data'.</p>
Threat Protection	Sends notifications about whether the threat protection service for Infoblox DNS Protection is functioning properly. For more information about the <code>ibPreviousState</code> (OID 3.1.1.1.2.9.0) and <code>ibCurrentState</code> (OID 3.1.1.1.2.10.0) tables, see <a href="#">SNMP MIB Hierarchy</a> .	<p>2016-04-18 22:48:08</p> <p>10.35.139.15 [UDP: [10.35.139.15]:54407-&gt;[10.120.21.204]]:</p> <p>DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (3926) 0:00:39.26</p> <p>SNMPv2-MIB::snmpTrapOID.0 = OID:</p> <p>IB-TRAP-MIB::ibStateChangeEvent</p> <p>IB-TRAP-MIB::ibNodeName.0 = STRING: "10.35.139.15"</p> <p>IB-TRAP-MIB::ibTrapSeverity.0 = INTEGER: info(2)</p> <p>IB-TRAP-MIB::ibObjectName.0 = STRING: threat_protection</p> <p>IB-TRAP-MIB::ibPreviousState.0 = INTEGER: threat-protection-service-inactive(93)</p> <p>IB-TRAP-MIB::ibCurrentState.0 = INTEGER: threat-protection-service-inactive(93)</p> <p>IB-TRAP-MIB::ibTrapDesc.0 = STRING: Threat Protection Service is inactive.</p>

## Testing the SNMP Configuration

After you configure SNMP on the appliance, you can do the following to test your SNMP configuration:

- From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox, and then select **Test SNMP** from the Toolbar.

The appliance sends a “test trap” string to the trap receiver and displays a confirmation message at the top of the screen if your SNMP configuration is properly set up. If your SNMP configuration is not complete or if it is invalid, the appliance displays an error message. You can check your configuration and try again.

The following is a sample test trap that the trap receiver can get:

```
2011-04-04 17:37:14 10.32.2.80 [UDP: [10.32.2.80]:49244->[10.32.2.80]]:
SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-MIB::snmpTrapOID
SNMPv2-MIB::sysName.0 = STRING: 'Test trap'
```

## Defining Interface for SNMP Traps

By default, the NIOS appliance sends SNMP traps through the MGMT interface to specific trap receivers, or through the LAN1 interface when the MGMT interface is disabled. However, you can choose an interface, instead of the default interface, to send SNMP traps to the trap receivers. For example, if you select LAN2 interface from the drop-down list, the traps are sent from the LAN2 interface instead of the default MGMT interface.

To choose an interface, other than MGMT or LAN1, for a Grid member and a standalone Grid:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *Grid\_member* checkbox.  
**Standalone:** From the **System** tab, select the **System Manager** tab, expand the Toolbar and click **System Properties** -> **Edit**.
2. In the editor, click **Toggle Advanced Mode**.
3. When the additional tabs appear, click the **Advanced** subtab of the **SNMP** tab.
4. Complete the following:
  - **Interface:** Select an interface from the drop-down list:
    - **ANY:** The default value is **ANY**. SNMP traps are sent through the MGMT interface, if it is enabled, when you select this option. If the interface is disabled, the traps are sent through the LAN1 interface.
    - **LAN1:** Select this option to send SNMP traps through the LAN1 interface.
    - **LAN2:** Select this option to send SNMP traps through the LAN2 interface.
    - **MGMT:** Select this option to send SNMP traps through the MGMT interface.
5. Click **Save & Close** to save the changes or **Cancel** to exit.

## SNMP MIB Hierarchy

In addition to implementing its own enterprise MIBs, Infoblox supports the standard MIBs defined in *RFC-1213, Management Information Base for Network Management of TCP/IP-based internets: MIB-II*.

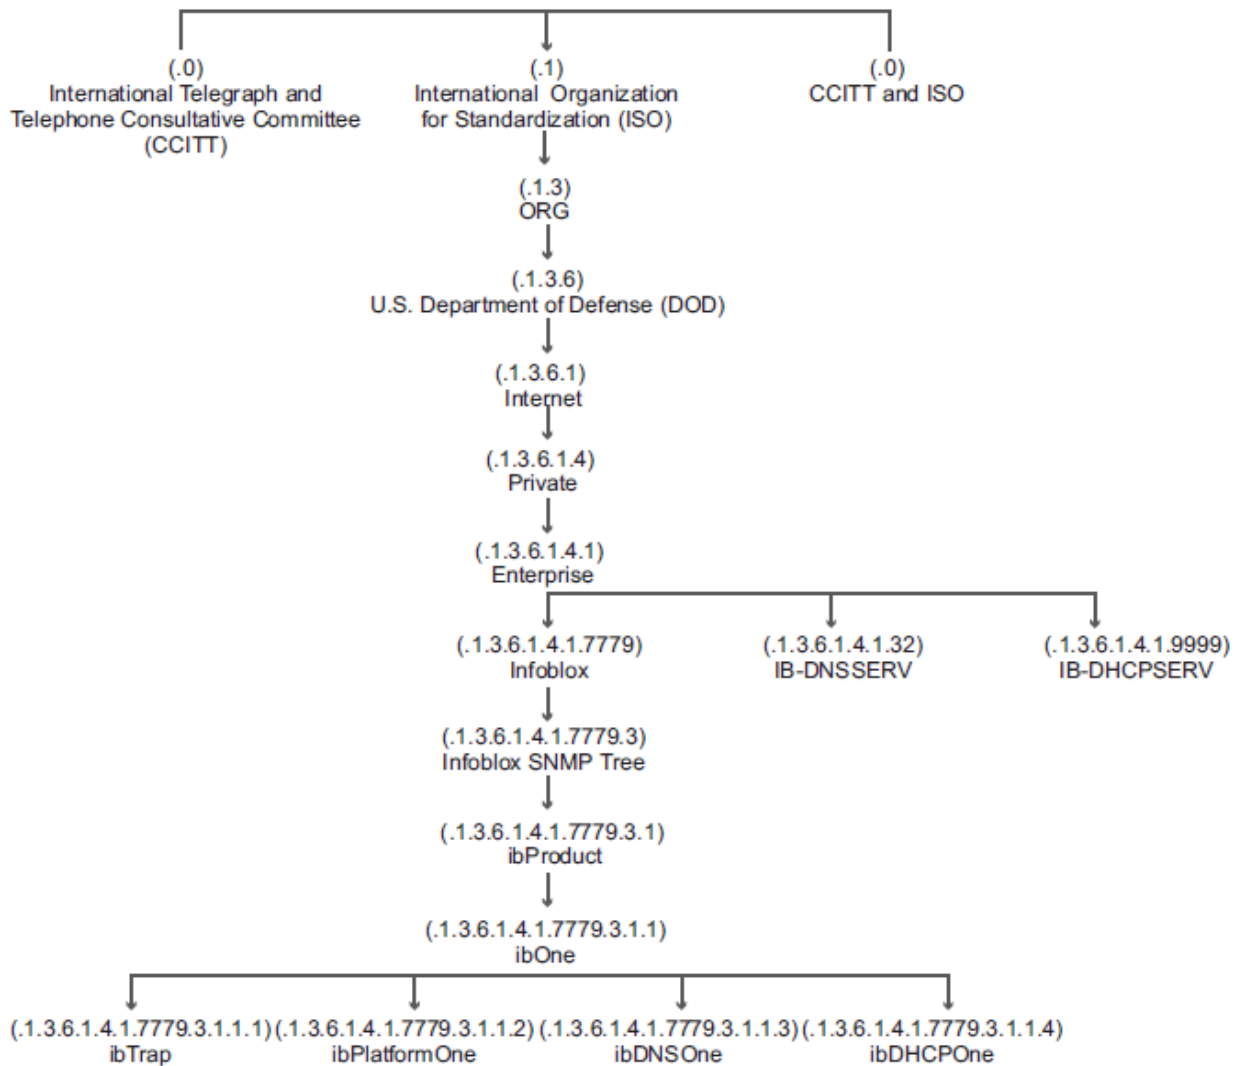
The Infoblox MIBs are part of a universal hierarchical structure, usually referred to as the MIB tree. The MIB tree has an unlabeled root with three subtrees. The below figure illustrates the branch of the MIB tree that leads to the Infoblox enterprise MIBs. Each object in the MIB tree has a label that consists of a textual description and an OID (object identifier). An OID is a unique dotted-decimal number that identifies the location of the object in the MIB tree. Note that all OIDs begin with a dot (.) to indicate the root of the MIB tree.

As shown in the figure below, Infoblox is a branch of the Enterprise subtree. IANA (Internet Assigned Numbers Authority) administers the Enterprise subtree, which is designated specifically for vendors who define their own MIBs. The IANA-assigned enterprise number of Infoblox is 7779; therefore, the OIDs of all Infoblox MIB objects begin with the prefix .1.3.6.1.4.1.7779. In addition, IB-DNSSERV and IB-DHCPSEV are branches of the Enterprise subtree as well.

The Infoblox SNMP subtree branches down through two levels, *ibProduct* and *ibOne*, to the Infoblox MIBs: *ibTrap*, *ibPlatformOne*, *ibDNSone*, and *ibDHCPone*. The *ibTrap* MIB defines the traps that NIOS appliances send, and the *ibPlatformOne*, *ibDNSone*, and *ibDHCPone* MIBs provide information about the appliance. For detailed information about these MIBs, see the *Infoblox MIBs* section.

### *MIB Hierarchy*





## MIB Objects

The Infoblox MIB objects were implemented according to the guidelines in RFCs 1155 and 2578. They specify two types of macros for defining MIB objects: OBJECT-TYPE and NOTIFICATION-TYPE. These macros contain clauses that describe the characteristics of an object, such as its syntax and its status. OBJECT-TYPE macros describe MIB objects, and NOTIFICATION-TYPE macros describe objects used in SNMP traps.

Each object in the ibPlatformOne, ibDNSOne, and ibDHCPOne MIBs contains the following clauses from the OBJECT-TYPE macro:

OBJECT-TYPE: Provides the administratively-assigned name of the object.

- SYNTAX: Identifies the data structure of the object, such as integers, counters, and octet strings.
- MAX-ACCESS: Identifies the type of access that a management station has to the object. All Infoblox MIB objects provide read-only access.
- STATUS: Identifies the status of the object. Values are current, obsolete, and deprecated.
- DESCRIPTION: Provides a textual description of the object.
- INDEX or AUGMENTS: An object that represents a conceptual row must have either an INDEX or AUGMENTS clause that defines a key for selecting a row in a table.

- **OID:** The dotted decimal object identifier that defines the location of the object in the universal MIB tree.

The ibTrap MIB defines the SNMP traps that a NIOS appliance can send. Each object in the ibTrap MIB contains the following clauses from the NOTIFICATION-TYPE macro:

NOTIFICATION-TYPE: Provides the administratively-assigned name of the object.

- **OBJECTS:** Provides an ordered list of MIB objects that are in the trap.
- **STATUS:** Identifies the status of the object. Values are current, obsolete, and deprecated.
- **DESCRIPTION:** Provides the notification information.

## System Object IDs

Infoblox uses the SNMP system object identifier sysObjectID to identify Infoblox appliances. The following table is a definition of sysObjectID from the SNMPv2 MIB, *Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)*:

<b>OBJECT-TYPE</b>	sysObjectID
SYNTAX	Object Identifier
MAX-ACCESS	read-only
STATUS	current
DESCRIPTION	"The vendor's authoritative identification of the network management subsystem contained in the entity. This value is allocated within the SMI enterprises subtree (1.3.6.1.4.1) and provides an easy and unambiguous means for determining 'what kind of box' is being managed. For example, if vendor 'Flintstones,Inc.' was assigned the subtree 1.3.6.1.4.1.424242, it could assign the identifier 1.3.6.1.4.1.424242.1.1 to its 'Fred Router'."

The following table lists the enterprise IDs and their corresponding Infoblox hardware platforms that an SNMP query can return when you request the sysObjectID value. Note that the IDs shown in the table do not include 1.3.6.1.4.1.7779.1. (the infobloxProducts prefix).

### *sysObjectID for Infoblox Hardware*

ID	Description	Definition
1000	ibDefault	Default environments, such as chroot
1002	ibCisco	Cisco servers
1003	ibvm	vNIOS appliances on VMware ESX or ESXi servers
1004	ibvnios	Virtual NIOS
1401	ib810	Trinzic 810 appliances
1402	ib820	Trinzic 820 appliances
1403	ib1410	Trinzic 1410 appliances

ID	Description	Definition
1404	ib1420	Trinzic 1420 appliances
1405	ib1400	Trinzic Reporting 1400 appliances
1406	ib800	Trinzic Reporting 800 appliances
1411	ib2200	Trinzic Reporting 2200 appliances
1412	ib2210	Trinzic 2210 appliances
1413	ib2220	Trinzic 2220 appliances
1423	ib4000	Infoblox-4000 appliances
1431	ib100	Trinzic 100 appliances
1444	nd4000	Network Insight 4000 appliances
1501	ib815	Trinzic 815 appliances
1502	ib825	Trinzic 825 appliances
1503	ib1415	Trinzic 1415 appliances
1504	ib1425	Trinzic 1425 appliances
1505	ib1405	Trinzic 1405 appliances
1506	ib805	Trinzic 805 appliances
1507	pt1405	Advanced 1405 appliances
1511	ib2205	Trinzic 2205 appliances
1512	ib2215	Trinzic 2215 appliances
1513	ib2225	Trinzic 2225 appliances
1514	pt2205	Advanced 2205 appliances
1521	ib4015	Trinzic 4015 appliances
1522	ib4035	Trinzic 4035 appliances

ID	Description	Definition
1523	ib4005	Trinzic 4005 appliances
1524	pt4005	Advanced 4005 appliances
1525	ib4025	Trinzic 4025 appliances
1541	nd805	Network Insight 805 appliances
1542	nd1405	Network Insight 1405 appliances
1543	nd2205	Network Insight 2205 appliances
1544	nd4005	Network Insight 4005 appliances

## Infoblox MIBs

You can configure a NIOS appliance as an SNMP-managed device so that an SNMP management station can send queries to the appliance and retrieve information from its MIBs. Perform the following tasks to access the Infoblox MIBs:

1. Configure a NIOS appliance to accept queries, as described in [Configuring SNMPv3 Users](#), see [Configuring SNMP](#).
2. Load the MIB files onto the management system. To obtain the latest Infoblox MIB files:
  - a. From the **Data Management** tab, select the **Grid** tab -> **Grid Manager** tab, and then select **Download** -> **SNMP MIBs** from the Toolbar.
  - b. In the Save As dialog box, navigate to a directory to which you want to save the MIBs.
  - c. Click Save.
1. Use a MIB browser or SNMP management application to query the objects in each MIB.

The NIOS appliance allows read-only access to the MIBs. This is equivalent to the Get and Get Next operations in SNMP.

## Loading the Infoblox MIBs

If you are using an SNMP manager toolkit with strict dependency checking, you must download the following Infoblox MIBs in the order they are listed:

1. IB-SMI-MIB.txt
2. IB-TRAP-MIB.txt
3. IB-PLATFORMONE-MIB.txt
4. IB-DNSONE-MIB.txt
5. IB-DHCPONE-MIB.txt
6. IB-DNSSERV-MIB.txt
7. IB-DHCPSERV-MIB.txt
8. IB-DHCPV6ONE-MIB.txt

In addition, if the SNMP manager toolkit you use requires a different MIB file naming convention, you can rename the MIB files accordingly.

## NET-SNMP MIBs

NIOS appliances support NET-SNMP (formerly UCD-SNMP), a collection of applications used to implement the SNMP protocol. The NET-SNMP MIBs provide the top-level infrastructure for the SNMP MIB tree. They define, among other

things, the objects in the SNMP traps that the agent sends when the SNMP engine starts and stops. For information about NET-SNMP and the MIB files distributed with NET-SNMP, refer to <http://net-snmp.sourceforge.net/>. For SNMP traps to function properly, you must download the following NET-SNMP MIBs directly from <http://net-snmp.sourceforge.net/docs/mibs>

- NET-SNMP-MIB
- UCD-SNMP-MIB



**Note**

Ensure that you save the MIBs as text files in the directory to which you save all the other MIB files.

### BGP4 MIB

Infoblox supports BGP4 (Border Gateway Protocol) for DNS anycast addressing. BGP is configured to send SNMP traps to neighboring routers, as defined in *RFC4273DefinitionsofManagedObjectsforBGP-4*. You must enable and configure the SNMP trap receiver on the Grid member for the member to send SNMP traps.

The BGP protocol service is configured to send SNMP queries about BGP runtime data. The information is returned using the following OIDs and definitions:

OID	Definition
1.3.6.1.2.1.15.900.1.1	Number of peers
1.3.6.1.2.1.15.900.1.2	Number of active peers
1.3.6.1.2.1.15.900.1.3	Number of AS path entries
1.3.6.1.2.1.15.900.1.4	Number of BGP community entries
1.3.6.1.2.1.15.900.1.5	Total number of prefixes

For each configured BGP peer (a, b, c, d), the information is returned using the following OIDs and definitions:

OID	Definition
1.3.6.1.2.1.15.900.1.9.a.b.c.d.1	IP address: same as a.b.c.d
1.3.6.1.2.1.15.900.1.9.a.b.c.d.2	State: 0=down, 1=up
1.3.6.1.2.1.15.900.1.9.a.b.c.d.3	ASN
1.3.6.1.2.1.15.900.1.9.a.b.c.d.4	Prefixes
1.3.6.1.2.1.15.900.1.9.a.b.c.d.5	Up/Down time

### ibTrap MIB

NIOS appliances send SNMP traps when events, internal process failures, or critical service failures occur. The ibTrap MIB defines the types of traps that a NIOS appliance sends and the value that each MIB object represents. The Infoblox

SNMP traps report objects which the ibTrap MIB defines. The below figure illustrates the ibTrap MIB structure. It provides the OID and textual description for each object.

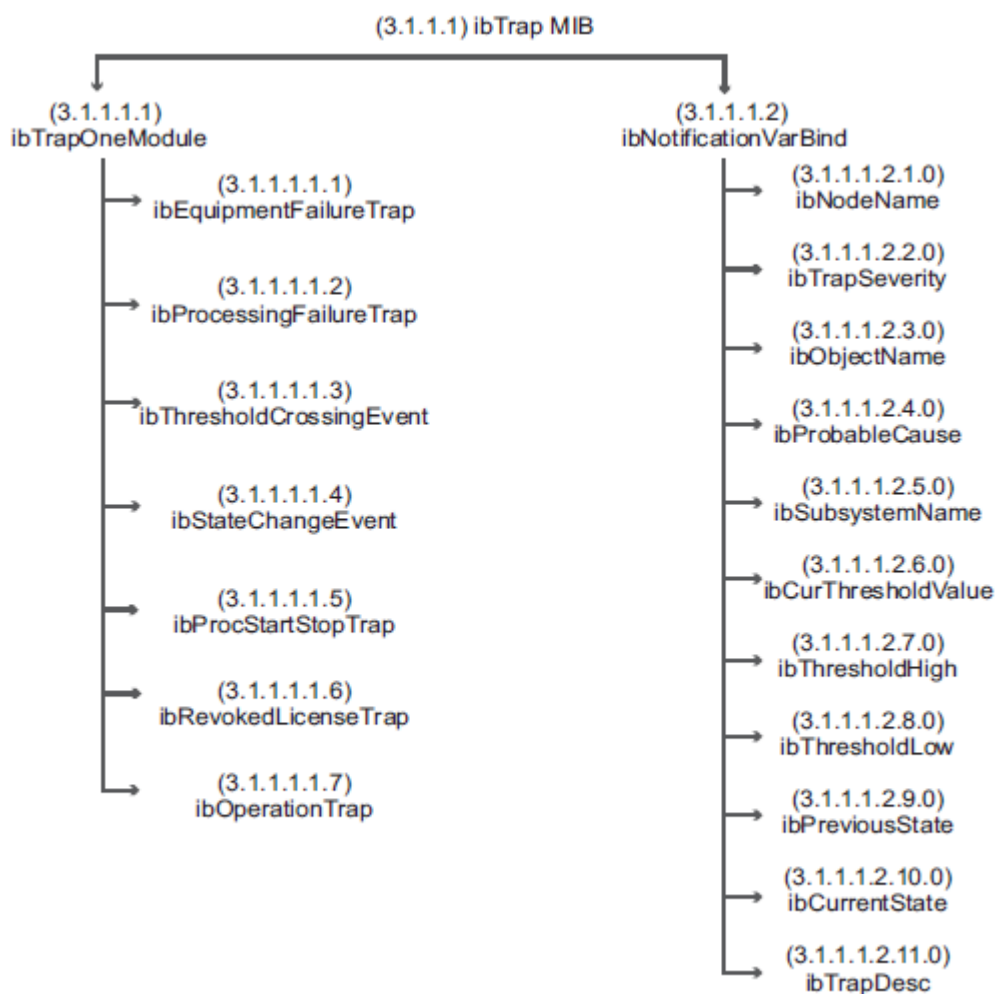


**Note**

OIDs shown in the illustrations and tables in this section do not include the prefix .1.3.6.1.4.1.7779.

The ibTrap MIB comprises two trees, ibTrapOneModule and ibNotificationVarBind. The ibTrapOneModule tree contains objects for the types of traps that a NIOS appliance sends. The ibNotificationVarBind tree contains objects that the Infoblox SNMP traps report. You cannot send queries for the objects in this MIB module. The objects are used only in the SNMP traps.

*ibTrapOne MIB Structure*



**Interpreting Infoblox SNMP Traps**

Depending on the SNMP management application your management system uses, the SNMP traps you receive may list the OIDs for all relevant MIB objects from both the ibTrapOneModule and ibNotificationVarBind trees. For OIDs that have string values, the trap lists the text. For OIDs that contain integers, you can use the tables in this section to find out the values. Some SNMP management applications list only the object names and their corresponding values in the SNMP

traps. Whether or not your SNMP management application lists OIDs, you can use the tables in this section to find out the corresponding value and definition for each MIB object.

The following is a sample trap a NIOS appliance sends:

```
418:Jan 31 18:52:26 (none) snmptrapd[6087]: 2008-01-31 18:52:26 10.35.1.156
[UDP:
[10.35.1.156]:32772]: DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks:
(1080)
0:00:10.80 SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-
SMI::enterprises.7779.3.1.1.1.4.0
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.1.0 = STRING: "10.35.1.156"
SNMPv2-SMI::enterprises.
7779.3.1.1.1.2.3.0 = STRING: "ntp_sync" SNMPv2-
SMI::enterprises.7779.3.1.1.1.2.9.0 =
INTEGER: 15 SNMPv2-SMI::enterprises.7779.3.1.1.1.2.10.0 = INTEGER: 16
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.11.0 = STRING: "The NTP service is
out of
synchronization."
```

The sample trap lists the OIDs and their corresponding values that can help you identify the cause of an event or problem. To identify the possible cause and recommended actions for the trap, use the `ibTrapDesc` tables.

You can interpret the sample trap as follows:

Using the `ibTrapOneModule` table, you find out OID 7779.3.1.1.1.4.0 represents an Object State Change trap. This trap includes the following objects: `ibNodeName`, `ibObjectName`, `ibPreviousState`, `ibCurrentState`, and `ibtrapDesc`. For each object, the trap displays the OID and its corresponding value. The following is how you can interpret the rest of the trap:

- `ibNodeName` (OID 7779.3.1.1.1.2.1.0)
  - Using the `ibNotificationVarBind` (OID 3.1.1.1.2) table, you find out OID 7779.3.1.1.1.2.1.0 represents the MIB object `ibNodeName`, which is the IP address of the appliance on which the trap occurred. Therefore, the statement `"7779.3.1.1.1.2.1.0 = STRING: "10.35.1.156" SNMPv2-SMI::enterprises."` tells you the IP address of the appliance on which the trap occurred.
- `ibObjectName` (OID 7779.3.1.1.1.2.3.0)
  - The statement `"7779.3.1.1.1.2.3.0 = STRING: "ntp_sync" SNMPv2-SMI::enterprises."` tells you the MIB object `ibObjectName`, which is the name of the object for which the trap was generated, has a value of "ntp\_sync" that indicates NTP synchronization issues.
- `ibPreviousState` (OID 7779.3.1.1.1.2.9.0)
  - The statement `"7779.3.1.1.1.2.9.0 = INTEGER: 15 SNMPv2-SMI::enterprises."` tells you the MIB object `ibPreviousState`, which indicates the previous state of the appliance, has a value of 15. Using the `ibPreviousState and ibCurrentState Values` table, you know that 15 represents "ntp-sync-up", which means the NTP server was up and running.
- `ibCurrentState` (OID 7779.3.1.1.1.2.10.0)



- The statement `"7779.3.1.1.1.2.10.0 = INTEGER: 16 SNMPv2-SMI::enterprises."` tells you the MIB object `ibCurrentState`, which indicates the current state of the appliance, has a value of 16. Using the *ibPreviousState and ibCurrentState Values* table, you know that 16 represents "ntp-sync-down", which means the NTP server is now out of sync.
- `ibTrapDesc` (OID 7779.3.1.1.1.2.11.0)
  - The last statement `"7779.3.1.1.1.2.11.0 = STRING: "The NTP service is out of synchronization."` states the description of the trap. Using the *Object State Change Traps* table for `ibTrapDesc`, you can find out the trap description and recommended actions for this problem.

### Types of Traps (OID 3.1.1.1.1)

`ibTrapOneModule` defines the types of traps that the NIOS appliance can send. There are five types of SNMP traps. The table below describe the types of traps and their objects in the `ibTrapOneModule` tree.



#### Note

Some SNMP traps for `ibThresholdCrossingEvent`, `ibStateChangeEvent`, `ibProcStartStopTrap`, and `ibRevokedLicenseTrap` do not have an associated `ibProbableCause`. The following table lists traps that provide `ibProbableCause` and those that do not have an `ibProbableCause` value.

### *ibTrapOneModule*

#### Trap Binding Variables (OID 3.1.1.1.2)

OID	Trap Type	MIB Object	Description
3.1.1.1.1.1	Equipment Failure	<code>ibEquipmentFailureTrap</code>	<p>The NIOS appliance generates this trap when a hardware failure occurs. This trap includes the following objects:</p> <ul style="list-style-type: none"> <li>• <code>ibNodeName</code></li> <li>• <code>ibTrapSevertiy</code></li> <li>• <code>ibObjectName</code> (equipment name)</li> <li>• <code>ibProbableCause</code></li> <li>• <code>ibTrapDesc</code></li> </ul>
3.1.1.1.1.2	Processing and Software Failure	<code>ibProcessingFailureTrap</code>	<p>The NIOS appliance generates this trap when a failure occurs in one of the software processes. This trap includes the following objects:</p> <ul style="list-style-type: none"> <li>• <code>ibNodeName</code></li> <li>• <code>ibTrapSeverity</code></li> <li>• <code>ibSubsystemName</code></li> <li>• <code>ibProbableCause</code></li> <li>• <code>ibTrapDesc</code></li> </ul>

3.1.1.1.1.3	Threshold Crossing	ibThresholdCrossingEvent	<p>The NIOS appliance generates this trap when any of the following events occur:</p> <ul style="list-style-type: none"> <li>• System memory or disk usage exceeds 90%.</li> <li>• CPU usage exceeds the trigger value for 15 seconds.</li> <li>• A problem occurs when the Grid Master replicates its database to its Grid members.</li> <li>• DHCP address usage crosses a watermark threshold.</li> <li>• The number or percentage of the DNS security alerts exceeds the thresholds of the DNS security alert triggers.</li> </ul> <p>This trap includes the following objects:</p> <ul style="list-style-type: none"> <li>• ibnodeName</li> <li>• ibTrapSeverity</li> <li>• ibObjectName (threshold name)</li> <li>• ibCurThresholdvalue</li> <li>• ibThresholdHigh</li> <li>• ibThresholdLow</li> <li>• ibTrapDesc</li> </ul>
3.1.1.1.1.4	Object State Change	ibStateChangeEvent	<p>The NIOS appliance generates this trap when there is a change in its state, such as:</p> <ul style="list-style-type: none"> <li>• The link to one of the configured ports goes down, and then goes back up again.</li> <li>• A failover occurs in an HA (high availability) pair configuration.</li> <li>• A member connects to the Grid Master.</li> <li>• An appliance in a Grid goes offline. This trap includes the following objects:</li> </ul> <ul style="list-style-type: none"> <li>• ibnodeName</li> <li>• ibTrapSeverity</li> <li>• ibObjectName</li> <li>• ibPreviousState</li> <li>• ibCurrentState</li> <li>• ibTrapDesc</li> </ul>

3.1.1.1.1.5	Process Started and Stopped	ibProcStartStopTrap	<p>The NIOS appliance generates this type of trap when any of the following events occur:</p> <ul style="list-style-type: none"> <li>• When you enable HTTP redirection.</li> <li>• When you change the HTTP access setting.</li> <li>• When you change the HTTP session time out setting.</li> <li>• When a failover occurs in an HA pair configuration.</li> </ul> <p>This trap includes the following objects:</p> <ul style="list-style-type: none"> <li>• <b>ibNodeName</b></li> <li>• <b>ibSubsystemName</b></li> <li>• <b>ibTrapDesc</b></li> </ul>
3.1.1.1.1.6		ibRevokedLicenseTrap	<p>The NIOS appliance generates this trap when a license is revoked.</p> <p>This trap includes the following objects:</p> <ul style="list-style-type: none"> <li>• <b>ibNodeName</b></li> <li>• <b>ibTrapSeverity</b></li> <li>• <b>ibSubsystemName</b></li> <li>• <b>ibTrapDesc</b></li> </ul>
3.1.1.1.1.7		ibOperationTrap	<p>The NIOS appliance generates this trap when a software operation is noteworthy.</p> <p>This trap includes the following objects:</p> <ul style="list-style-type: none"> <li>• <b>ibNodeName</b></li> <li>• <b>ibTrapSeverity</b></li> <li>• <b>ibSubsystemName</b></li> <li>• <b>ibProbableCause</b></li> <li>• <b>ibTrapDesc</b></li> </ul>

Each SNMP trap contains information about the event or the problem. The Infoblox SNMP traps include MIB objects and their corresponding values from the `ibNotificationVarBind` module. The following table describes the objects in the `ibNotificationVarBind` module.

The OIDs shown in the following table do not include the prefix ".1.3.6.1.4.1.7779.":

*ibNotificationVarBind (OID 3.1.1.1.2)*

OID	MIB Object (Type)	Description
3.1.1.1.2.1.0	ibNodeName (DisplayString)	The IP address of the appliance on which the trap occurs. This may or may not be the same as the appliance that sends the trap. This object is used in all types of traps.

3.1.1.1.2.2.0	ibTrapSeverity (Integer)	The severity of the trap. There are five levels of severity. See the Trap Severity table (OID 3.1.1.1.2.2.0) for details below.
3.1.1.1.2.3.0	ibObjectName (DisplayString)	The name of the object for which the trap was generated. This is used in the Equipment Failure traps, Threshold Crossing Event traps, and the Object State Change traps. The following shows what this object represents depending on the type of traps: <ul style="list-style-type: none"> <li>• Equipment Failure traps: The equipment name.</li> <li>• Threshold Crossing Event traps: The object name of the trap.</li> <li>• State Change traps: The object that changes state.</li> </ul>
3.1.1.1.2.4.0	ibProbableCause (Integer)	The probable cause of the trap.
3.1.1.1.2.5.0	ibSubsystemName (DisplayString)	The subsystem for which the trap was generated, such as NTP or SNMP. This object is used in the Processing and Software Failure traps and the Process Start and Stop traps.
3.1.1.1.2.6.0	ibCurThresholdValue (Integer)	The current value of the threshold counter. This object is used in the Threshold Crossing traps.
3.1.1.1.2.7.0	ibThresholdHigh (Integer)	This object is used in Threshold Crossing traps. For CPU usage, this is the trigger value of the SNMP trap. For DHCP address usage, this is the value of the high watermark. This only applies when the appliance sends a trap to indicate that DHCP address usage is above the configured high watermark value for a DHCP address range.
3.1.1.1.2.8.0	ibThresholdLow (Integer)	This object is used in Threshold Crossing traps. For CPU usage, this is the reset value of the SNMP trap. For DHCP address usage, this is the value for the low watermark. This only applies when the appliance sends a trap to indicate that DHCP address usage went below the configured low watermark value for a DHCP address range.
3.1.1.1.2.9.0	ibPreviousState (Integer)	The previous state of the appliance. This object is used in the Object State Change traps. See the <i>ibPreviousState</i> (OID 3.1.1.1.2.9.0) and <i>ibCurrentState</i> (OID 3.1.1.1.2.10.0) section for definitions of each value.
3.1.1.1.2.10.0	ibCurrentState (Integer)	The current state of the appliance. This object is used in the Object State Change traps. See the <i>ibPreviousState</i> (OID 3.1.1.1.2.9.0) and <i>ibCurrentState</i> (OID 3.1.1.1.2.10.0) section for definitions of each value.
3.1.1.1.2.11.0	ibTrapDesc (DisplayString)	The description of the trap. This object is used in all types of traps. See the <i>ibTrapDesc</i> (OID 3.1.1.1.2.11.0) section for the description, possible cause, and recommended actions for each Infoblox SNMP trap.

### Trap Severity (OID 3.1.1.1.2.2.0)

The object `ibTrapSeverity` defines the severity level for each Infoblox SNMP trap. There are five levels of severity:

Value	Description
1	Indetermined [sic]
2	Informational: Event that requires no further action.
3	Minor: Event that does not require user intervention.
4	Major: Event that requires user intervention and assistance from Infoblox Technical Support.
5	Critical: Problem that affects services and system operations, and requires assistance from Infoblox Technical Support.

### `ibProbableCause` Values (OID 3.1.1.1.2.4.0)

The following table lists the values that are associated with the object `ibProbableCause` (OID 3.1.1.1.2.4.0). These values provide information about the events such as hardware, software for process failures, that trigger SNMP traps.

#### *ibProbableCause Values*

Value	OID 3.1.1.1.2.4.0 <code>ibProbableCause</code>	Equipment, Software, or Process Failure Traps
0	<code>ibClear</code>	SNMP Trap is cleared.
1	<code>ibUnknown</code>	An unknown failure has occurred.
2	<code>ibPrimaryDiskFailure</code>	A primary drive failure has occurred.
3	<code>ibFanFailure-old</code>	Unused.
4	<code>ibPowerSupplyFailure</code>	A power supply failure has occurred.
5	<code>ibDBFailure</code>	A database daemon monitoring failure has occurred.
6	<code>ibApacheSoftwareFailure</code>	An apache software failure has occurred.
7	<code>ibSerialConsoleFailure</code>	An Infoblox serial console software failure has occurred.
11	<code>ibControldSoftwareFailure</code>	A controld failure has occurred.
12	<code>ibUpgradeFailure</code>	A system upgrade failure has occurred.
13	<code>ibSNMPDFailure</code>	SNMP Server failure has occurred.

15	ibSSHDSOFTWAREFAILURE	An SSH daemon failure has occurred.
16	ibNTPDSOFTWAREFAILURE	An NTP daemon failure has occurred.
17	ibClusterdSOFTWAREFAILURE	A cluster daemon failure has occurred.
18	ibLCDSOFTWAREFAILURE	An LCD daemon failure has occurred.
19	ibDHCPdSOFTWAREFAILURE	A DHCP daemon monitoring failure has occurred.
20	ibNamedSOFTWAREFAILURE	A named daemon monitoring failure has occurred.
21	ibAuthServerGroupDown	NAC Authentication server group is down.
22	ibAuthServerGroupUp	NAC Authentication server group is up.
24	ibNTLMSOFTWAREFAILURE	An NTLM monitoring failure has occurred.
25	ibNetBIOSDaemonFailure	A NetBIOS daemon monitoring failure has occurred.
26	ibWindowBindDaemonFailure	An NT domain service monitoring failure has occurred.
27	ibTFTPDSoftwareFailure	A TFTP daemon failure has occurred.
28	ibUNUSED28	Unused.
29	ibBackupSoftwareFailure	Backup failed.
30	ibBackupDatabaseSoftwareFailure	Database backup failed.
31	ibBackupModuleSoftwareFailure	Module backup failed.
32	ibBackupSizeSoftwareFailure	File size exceeded the quota. Backup failed.
33	ibBackupLockSoftwareFailure	Another backup is in progress. Backup will not be performed.
34	ibHTTPFileDistSoftwareFailure	An HTTP file distribution daemon failure has occurred.
35	ibOSPFSOFTWAREFAILURE	An OSPF routing daemon failure has occurred.
36	ibAuthDHCPNamedSoftwareFailure	An auth named server failure has occurred.
37	ibFan1Failure	Fan 1 failure has occurred.
38	ibFan2Failure	Fan 2 failure has occurred.

39	ibFan3Failure	Fan 3 failure has occurred.
40	ibFan1OK	Fan 1 is OK.
41	ibFan2OK	Fan 2 is OK.
42	ibFan3OK	Fan 3 is OK.
44	ibFTPDSoftwareFailure	An FTPD daemon failure has occurred.
45	ibBloxtoolsSoftwareFailure	A Bloxtools service failure has occurred.
46	ibPowerSupplyOK	The power supply is OK.
47	ibWebUISoftwareFailure	A WebUI software failure has occurred.
48	ibUNUSED48	Unused.
49	ibADAgentSyncFailure	An AD agent client synchronizing domain data failure has occurred.
50	ibIFMAPSoftwareFailure	An IF-MAP server failure has occurred.
51	ibCaptivePortalSoftwareFailure	A Captive Portal service failure has occurred.
52	ibDuplicateIPAddressFailure	A Duplicate IP Address has been detected.
53	ibBGPSoftwareFailure	An BGP routing daemon failure has occurred.
54	ibRevokedLicense	A license has been revoked.
58	ibGUILoginFailure	An admin failed to log in to the GUI.
59	ibSerialConsoleLoginFailure	An admin failed to log in to the serial console.
60	ibSystemReboot	A system reboot was initiated.
61	ibSystemRestart	A system restart was initiated.
62	ibZoneTransferFailure	A zone transfer failure occurred.
63	ibDHCPLeaseConflict	DHCP address conflicts with an existing lease.
64	ibDHCPAddressConflict	DHCP address conflicts with an existing fixed address.
65	ibDHCPRangeConflict	DHCP address conflicts with an existing range.



66	ibDHCPHostConflict	DHCP address conflicts with an existing host.
67	ibSyslogFailure	A syslog daemon failure occurred.
68	ibPowerSupply1Failure	Power supply 1 failure has occurred.
69	ibPowerSupply2Failure	Power supply 2 failure has occurred.
70	ibPowerSupply1OK	Power supply 1 is OK.
71	ibPowerSupply2OK	Power supply 2 is OK.
72	ibReportingTaskSwFailure	A reporting task monitoring failure has occurred.
73	ibReportingDbBackupFailure	A reporting db backup/restore operation failure has occurred.
74	ibFan4Failure	Fan 4 failure has occurred.
75	ibFan5Failure	Fan 5 failure has occurred.
76	ibFan6Failure	Fan 6 failure has occurred.
77	ibFan7Failure	Fan 7 failure has occurred.
78	ibFan8Failure	Fan 8 failure has occurred.
79	ibFan4OK	Fan 4 is OK.
80	ibFan5OK	Fan 5 is OK.
81	ibFan6OK	Fan 6 is OK.
82	ibFan7OK	Fan 7 is OK.
83	ibFan8OK	Fan 8 is OK.
84	ibOSPF6SoftwareFailure	An OSPF6 routing daemon failure has occurred.
85	ibOCSPResponderFailure	OCSP responder failed.
86	ibReportingAlertTriggered	A reporting alert is triggered.
87	ibCapturedQueriesUploadFailure	Upload for captured DNS queries failed.
88	ibLDAPServerFailure	The LDAP server failed.

89	ibRIRWIPRegistrationFailure	RIR SWIP registration failed.
90	ibPowerSupply1Removed	Power supply 1 has been removed.
91	ibPowerSupply2Removed	Power supply 2 has been removed.
92	ibIPMISensorErrorDetected	Error detected on the sensor for the IPMI port (used for LOM)
93	ibDiscoveryConsolidatorTaskSwFailure	Discovery service on the Consolidator failed.
94	ibDiscoveryCollectorTaskSwFailure	Discovery service on probes failed.
95	ibDiscoveryBackupSwFailure	Discovery backup service failed.
96	ibThreatProtectionAutoDownloadFailure	Automatic download of threat protection rule failed.
97	ibThreatProtectionPublishFailure	Threat protection rule publish failed.
98	ibPassiveHANodeARPCConnectivityFailure	HA node failed to connect to local router.
99	ibPassiveHANodeARPCConnectivitySuccess	HA node successfully connects to local router.
100	ibDNSIntegrityCheckConnectionFailed	Connection between Grid member and Grid Master failed.
101	ibDNSIntegrityCheckPrimaryServersFailed	DNS data (NS RRset) check for Grid primaries failed.
102	ibDNSIntegrityCheckNameserversFailed	DNS data (NS RRset) check for name servers failed.
103	ibCloudAPIFailure	Cloud API service failed.
104	ibRpzRefreshFailure	An RPZ refresh failure has occurred.
105	ibUnboundSoftwareFailure	Unbound software failure has occurred.
106	ibAnalyticsAutoDownloadFailure	Automatic download of analytic module set failed.
107	ibDnsHealthCheckFailed	The DNS health check has failed.
108	ibDnsHealthCheckSucceed	The DNS health check is successful.
109	ibBFDSsoftwareFailure	BFD has failed to detect failure in the bidirectional path between two interfaces.
110	ibOutboundWorkerFailed	An Outbound worker has failed.
3001	ibRAIDIsOptimal	The system's RAID array is now running in an optimal state.

3002	ibRAIDIsDegraded	The system's RAID array is in a degraded state.
3003	ibRAIDIsRebuilding	The system's RAID array is rebuilding.
3004	ibRAIDStatusUnknown	Unable to retrieve RAID array state!
3005	ibRAIDBatteryIsOK	The system's RAID battery is OK.
3006	ibRAIDBatteryFailed	A RAID battery failure has occurred.
3007	ibRAIDOptimalMismatch	The system's RAID array is now running in an optimal state (Mismatched disk(s) found).
3008	ibRAIDDegradedMismatch	The system's RAID array is in a degraded state (Mismatched disk(s) found).
3009	ibRAIDRebuildingMismatch	The system's RAID array is rebuilding (Mismatched disk(s) found).
3010	ibRAIDBatteryWeak	Please replace the system's RAID battery soon.
3011	ibRAIDIsDegradedDisk1	The system's RAID array is in a degraded state. RAID Disk1 is EMPTY.
3012	ibRAIDIsDegradedDisk2	The system's RAID array is in a degraded state. RAID Disk2 is EMPTY.
3013	ibRAIDIsDegradedDisk3	The system's RAID array is in a degraded state. RAID Disk3 is EMPTY.
3014	ibRAIDIsDegradedDisk4	The system's RAID array is in a degraded state. RAID Disk4 is EMPTY.
3015	ibRAIDIsDegradedDisk5	The system's RAID array is in a degraded state. RAID Disk5 is EMPTY.
3016	ibRAIDIsDegradedDisk6	The system's RAID array is in a degraded state. RAID Disk6 is EMPTY.
3017	ibRAIDIsDegradedDisk7	The system's RAID array is in a degraded state. RAID Disk7 is EMPTY.
3018	ibRAIDIsDegradedDisk8	The system's RAID array is in a degraded state. RAID Disk8 is EMPTY.
3019	ibRAIDIsRebuildingDisk1	The system's RAID array is rebuilding. RAID Disk1 is OFFLINE.
3020	ibRAIDIsRebuildingDisk2	The system's RAID array is rebuilding. RAID Disk2 is OFFLINE.
3021	ibRAIDIsRebuildingDisk3	The system's RAID array is rebuilding. RAID Disk3 is OFFLINE.
3022	ibRAIDIsRebuildingDisk4	The system's RAID array is rebuilding. RAID Disk4 is OFFLINE.
3023	ibRAIDIsRebuildingDisk5	The system's RAID array is rebuilding. RAID Disk5 is OFFLINE.
3024	ibRAIDIsRebuildingDisk6	The system's RAID array is rebuilding. RAID Disk6 is OFFLINE.

3025	ibRAIDIsRebuildingDisk7	The system's RAID array is rebuilding. RAID Disk7 is OFFLINE.
3026	ibRAIDIsRebuildingDisk8	The system's RAID array is rebuilding. RAID Disk8 is OFFLINE.
3027	ibRAIDDegradedMismatchDisk1	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk1 is EMPTY.
3028	ibRAIDDegradedMismatchDisk2	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk2 is EMPTY.
3029	ibRAIDDegradedMismatchDisk3	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk3 is EMPTY.
3030	ibRAIDDegradedMismatchDisk4	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk4 is EMPTY.
3031	ibRAIDDegradedMismatchDisk5	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk5 is EMPTY.
3032	ibRAIDDegradedMismatchDisk6	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk6 is EMPTY.
3033	ibRAIDDegradedMismatchDisk7	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk7 is EMPTY.
3034	ibRAIDDegradedMismatchDisk8	The system's RAID array is in a degraded state (Mismatched disk(s) found). RAID Disk8 is EMPTY.
3035	ibRAIDRebuildingMismatchDisk1	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk1 is OFFLINE.
3036	ibRAIDRebuildingMismatchDisk2	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk2 is OFFLINE.
3037	ibRAIDRebuildingMismatchDisk3	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk3 is OFFLINE.
3038	ibRAIDRebuildingMismatchDisk4	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk4 is OFFLINE.
3039	ibRAIDRebuildingMismatchDisk5	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk5 is OFFLINE.
3040	ibRAIDRebuildingMismatchDisk6	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk6 is OFFLINE.
3041	ibRAIDRebuildingMismatchDisk7	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk7 is OFFLINE.
3042	ibRAIDRebuildingMismatchDisk8	The system's RAID array is rebuilding (Mismatched disk(s) found). RAID Disk8 is OFFLINE.

3043	ibRAIDIsInoperable	The system's RAID array is inoperable.
3044	ibRAIDPredictedDegraded	The system's RAID array is in a predicted degraded state.
3045	ibRAIDPredictedDegradedDisk1	The system's RAID array is in a predicted degraded state. RAID Disk1 is in such predicted degraded state.
3046	ibRAIDPredictedDegradedDisk2	The system's RAID array is in a predicted degraded state. RAID Disk2 is in such predicted degraded state.
3047	ibRAIDPredictedDegradedDisk3	The system's RAID array is in a predicted degraded state. RAID Disk3 is in such predicted degraded state.
3048	ibRAIDPredictedDegradedDisk4	The system's RAID array is in a predicted degraded state. RAID Disk4 is in such predicted degraded state.
3049	ibRAIDPredictedDegradedDisk5	The system's RAID array is in a predicted degraded state. RAID Disk5 is in such predicted degraded state.
3050	ibRAIDPredictedDegradedDisk6	The system's RAID array is in a predicted degraded state. RAID Disk6 is in such predicted degraded state.
3051	ibRAIDPredictedDegradedDisk7	The system's RAID array is in a predicted degraded state. RAID Disk7 is in such predicted degraded state.
3052	ibRAIDPredictedDegradedDisk8	The system's RAID array is in a predicted degraded state. RAID Disk8 is in such predicted degraded state.
4001	ibDisconnectGridAttachFailed	A disconnected Grid failed to attach to the Master Grid.
4002	ibDisconnectedGridDetachFailed	A disconnected Grid failed to detach from the Master Grid.
4003	ibDisconnectedGridDetachFailedSubgrid Offline	An offline Grid failed to detach from the Master Grid.
4004	ibDisconnectedGridSnapshotFailed	The snapshot operation failed on a disconnected Grid.
4005	ibDnssecAutomaticKSKRolloverApproching	An automatic rollover of the KSK will be performed in seven days.
4006	ibDnssecManualKSKRolloverDueApproching	A manual KSK rollover is due within the next seven days.
4007	ibDnssecAutomaticKSKRolloverDone	The KSK has been automatically rolled over.
4008	ibDnssecManualKSKRolloverDone	The KSK has been manually rolled over.
4009	ibDnssecKSKRolloverOverdue	The KSK rollover is overdue.
4010	ibPortDiscoveryConflict	A port detection conflict has been detected.

4011	ibDeviceDiscoveryConflict	A device detection conflict has been detected.
4012	ibDeviceUnmanaged	New unmanaged devices/networks were found during network discovery process.
4013	ibAnalyticsDnstUpdate	The DNS tunneling detection has occurred.
4014	ibSyslogBackupDone	The syslog backup process is successful.
4015	ibSyslogBackupFailed	The syslog backup process failed.

#### ibSubsystemName Values (OID 3.1.1.1.2.5.0)

The below table lists the values that are associated with the object `ibSubsystemName` (OID 3.1.1.1.2.5.0). These values provide information about the subsystems that trigger the traps.

#### *ibSubsystemName Values*

Value	OID 3.1.1.1.2.5.0 <code>ibSubsystemName</code>
0	Uses the original <code>ibObjectName</code> and <code>ibSubsystemName</code> when the trap is cleared. The process failure trap is appended to the CLEAR trap descriptions.
1	N/A
2	N/A
3	N/A
4	N/A
5	Db_jnlld
6	httpd
7	serial_console
11	controld
12	N/A
13	Snmpd
15	Sshd
16	Ntpd
17	Clusterd

18	Lcd
19	Dhcpd
20	Named
24	NTLM
25	Netbiosd
26	Winbindd
27	Tftpd
29	N/A
30	db_dump
31	N/A
32	Scheduled_backups
33	N/A
34	HTTPd
35	OSPF
36	AuthDhcpNamed
44	ftpd
45	bloxtools
47	webui
50	ifmap
51	captive_portal
53	BGP
58	GUI_Login
72	Reporting_task



84	OSPF6
95	Discovery_backups
105	Unbounds
-	HA_Replication
-	HSM Group Status
-	Cluster
-	Cluster_Send_Queue
-	Cluster_Recv_Queue
-	RPZHitRate
-	Authentication Server Group

ibPreviousState (OID 3.1.1.1.2.9.0) and ibCurrentState (OID 3.1.1.1.2.10.0)

The ibPreviousState object indicates the state of the appliance before the event triggered the trap. The ibCurrentState object indicates the current state of the appliance. The following table shows the message and description for each state:

*ibPreviousState and ibCurrentState Values*

Value	Description	Definition
1	ha-active	The HA pair is in ACTIVE state.
2	ha-passive	The HA pair is in PASSIVE state.
3	ha-initial	The HA pair is in INITIAL state.
4	Grid-connected	The Grid member is connected to the Grid Master.
5	Grid-disconnected	The Grid member is not connected to the Grid Master.
12	service-down	The service is down.
13	ha-replication-online	The HA pair replication is online.
14	ha-replication-offline	The HA pair replication is offline.

15	ntp-syn-up	The NTP server is synchronizing.
16	ntp-syn-down	The NTP server is out of synchronization.
17	ms-server-up	Microsoft server is up.
18	ms-server-down	Microsoft server is down.
19	ms-service-up	Microsoft service connection is active.
20	ms-service-down	Microsoft service connection is inactive.
21	nac-server-group-down	NAC Authentication server group is down.
22	nac-server-group-up	NAC Authentication server group is up.
23	mgm-service-up	MGM service is active.
24	mgm-service-down	MGM service is inactive.
25	ha-active-active	HA Pair is in Dual Active state.
26	ftp-service-working	FTP service is working.
27	ftp-service-failed	FTP service failed.
28	ftp-service-inactive	FTP service is inactive.
29	tftp-service-working	TFTP service is working.
30	tftp-service-failed	TFTP service failed.
31	tftp-service-inactive	TFTP service is inactive.
32	dns-service-working	DNS service is working.
33	dns-service-failed	DNS service failed. Review the syslog file.
34	dns-service-inactive	DNS service is inactive. Check if an administrator disabled the service.
35	ntp-service-working	NTP service is working.
36	ntp-service-failed	NTP service failed.
37	ntp-service-inactive	NTP service is inactive.

38	http-file-dist-service-working	HTTP File Dist service is working.
39	http-file-dist-service-failed	HTTP File Dist service failed.
40	http-file-dist-service-inactive	HTTP File Dist service is inactive.
41	bloxtools-service-working	bloxTools service is working.
42	bloxtools-service-warning	bloxTools service is in warning state.
43	bloxtools-service-failed	bloxTools service failed.
44	bloxtools-service-inactive	bloxTools service is inactive.
45	dhcp-service-working	DHCP service is working. No action required.
46	dhcp-service-warning	DHCP service is in warning state. Review the syslog file.
47	dhcp-service-failed	DHCP service failed. Review the syslog file.
48	dhcp-service-inactive	DHCP service is inactive. Check if an administrator disabled the service.
49	captive-portal-service-working	Captive portal service is working.
50	captive-portal-service-failed	Captive portal service failed.
51	captive-portal-service-inactive	Captive portal service inactive.
52	ifmap-service-working	IF-MAP service is working.
53	ifmap-service-failed	IF-MAP service failed.
54	ifmap-service-inactive	IF-MAP service inactive.
56	hsm-group-down	HSM operation failed.
57	hsm-group-up	HSM operation succeeded.
59	reporting-service-working	Reporting service is working.
60	reporting-service-failed	Reporting service failed.
61	reporting-service-inactive	Reporting service inactive.
62	dns-cache-acceleration-working	DNS cache acceleration is working.

63	dns-cache-acceleration-failed	DNS cache acceleration failed.
64	dns-cache-acceleration-inactive	DNS cache acceleration is inactive.
65	ocsp-responders-ok	All OCSP responders are available.
66	ocsp-responder-failed	At least one OCSP responder has become unavailable.
67	ocsp-responders-unavailable	All OCSP responders are out of service.
68	cas-inactive	Certificate authentication service is inactive.
69	subgrid-attached	In a Multi-Grid configuration, sub Grid is attached to the Master Grid.
70	subgrid-detached	In a Multi-Grid configuration, sub Grid is detached from the Master Grid.
71	snapshot-disabled	In a Multi-Grid configuration, the snapshot of the Grid's current state is disabled.
72	snapshot-enabled	In a Multi-Grid configuration, the snapshot of the Grid's current state is enabled.
73	lb-device-up	Load balancer device is working.
74	lb-device-down	Load balancer device is down.
76	ldap-servers-ok	All LDAP servers are available.
77	ldap-server-failure	At least one LDAP server is unavailable.
78	ldap-servers-unavailable	All LDAP servers are unavailable.
79	ldap-service-inactive	LDAP service is inactive.
80	sgm-state-online	SGM is online.
81	sgm-state-offline	SGM is offline.
82	discovery-consolidator-service-working	Service on discovery the consolidator for Network Insight is working.
83	discovery-consolidator-service-warning	Service on discovery the consolidator for Network Insight is in warning state.
84	discovery-consolidator-service-failed	Service on discovery the consolidator for Network Insight failed.
85	discovery-consolidator-service-inactive	Service on discovery the consolidator for Network Insight is inactive.

86	discovery-collector-service-working	Service on discovery probes for Network Insight is working.
87	discovery-collector-service-warning	Service on discovery probes for Network Insight is in warning state.
88	discovery-collector-service-failed	Service on discovery probes for Network Insight failed.
89	discovery-collector-service-inactive	Service on discovery probes for Network Insight is inactive.
90	threat-protection-service-working	Threat protection service for Infoblox DNS Protection is working.
91	threat-protection-service-warning	Threat protection service for Infoblox DNS Protection is in warning state
92	threat-protection-service-failed	Threat protection service for Infoblox DNS Protection failed.
93	threat-protection-service-inactive	Threat protection service for Infoblox DNS Protection is inactive.
94	mgm-external-storage-disabled	MGM external storage is disabled.
95	dns-integrity-check-failed	DNS Integrity check for authoritative zones failed.
96	dns-integrity-check-working	DNS Integrity check for authoritative zones is working.
97	dns-integrity-check-severity-indetermined	DNS Integrity check severity is None.
98	dns-integrity-check-severity-normal	DNS Integrity check severity is NORMAL.
99	dns-integrity-check-severity-informational	DNS Integrity check severity is INFORMATIONAL.
100	dns-integrity-check-severity-warning	DNS Integrity check severity is WARNING.
101	dns-integrity-check-severity-severe	DNS Integrity check severity is SEVERE.
102	dns-integrity-check-severity-critical	DNS Integrity check severity is CRITICAL.
103	cloud-api-service-working	Cloud API service is working.
104	cloud-api-service-failed	Cloud API service failed.

105	cloud-api-service-inactive	Could API service is inactive.
106	dns-service-dtc-failed	DNS sub-feature DNS Traffic Control failed.
107	raid-status-is-unavailable	RAID hardware status is unavailable.
108	raid-status-is-available	RAID hardware status is available again.
109	ipmi-device-warning	IPMI device is not available.
110	ipmi-device-unavailable	IPMI device repeatedly is not available.
111	ipmi-device-available	IPMI device is available again.
112	rpz-refresh-working	RPZ refresh is working.
113	rpz-refresh-failed	RPZ refresh failed.
114	rpz-refresh-unknown	RPZ refresh unknown.
115	dns-attack-active	DNS attack is active.
116	dns-attack-inactive	DNS attack is inactive.
117	high-rpz-hit-rate	RPZ hit rate exceeds the normal value.
118	normal-rpz-hit-rate	RPZ hit rate has returned to normal value.
119	taxii-working	TAXII server is working.
120	taxii-warning	TAXII server is warning.
121	taxii-failed	TAXII server is failed.
122	taxii-inactive	TAXII server is inactive.
123	cisco-ise-server-up	Cisco ISE server is connected.
124	cisco-ise-server-down	Cisco ISE server is not connected.
125	threat-analytics-service-inactive	Threat Analytics service is inactive.
126	threat-analytics-service-working	Threat Analytics service is working.
127	threat-analytics-service-warning	Threat Analytics service is in warning state.

128	threat-analytics-service-failed	Threat Analytics Service is failed.
129	outbound-service-manager-started	OUTBOUND Service Manager is working
130	outbound-service-manager-failed	OUTBOUND Service Manager is failed
131	outbound-service-manager-stopped	OUTBOUND Service Manager is inactive
133	imc-servers-ok	Subscriber Collection Service is working.
134	imc-server-failure	Subscriber Collection Service failed.
135	imc-servers-unavailable	All NAS gateways are unavailable.
136	imc-service-inactive	Subscriber Collection Service is inactive.
137	imc-zvelo-service-unavailable	Initial category information data is downloading.
138	imc-server-initializing	Initial Subscriber Collection service interim interval.

The Grid member running Subscriber Collection Service has a default state of 136 (imc-service-inactive). When Subscriber Collection Service is enabled, it goes to the 138 (imc-server-initializing) state. That is, the service will be enabled but not fully populated until the interim interval expires. Once the interim interval expires, it transitions to the 133 (imc-servers-ok) state unless parental control is enabled. If parental control is enabled, the state changes to 137 (imc-zvelo-service-unavailable) until the category data is downloaded. After this it goes to state 133 (imc-servers-ok). If the collector service is not running, it transitions to state 134 (imc-server-failure).

Trap 136 implies that subscriber collection is not active. Trap 134 implies that the collector is not running. If the monitor is disabled, then trap 136 (imc-service-inactive) is triggered.

 Note

- When downloading the category data, it is around 2.8 GB when compressed and around 6 GB when uncompressed. This data size remains unchanged even for incremental downloads. Ensure that the proxy settings are modified accordingly, before the download.
- The data from the earlier service provider is deleted after the upgrade and the categorization will not work until the categorization DB is downloaded.

### ibTrapDesc (OID 3.1.1.1.2.11.0)

The ibTrapDesc object lists the trap messages of all Infoblox SNMP traps. This section lists all the SNMP traps by their trap types. Each trap table describes the trap message, severity, cause, and recommended actions.

 Note

Contact Infoblox Technical Support for assistance when the recommended actions do not resolve the problems.

### IEquipment Failure Traps



ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity OID 3.1.1.1.2.2	Description/Cause	Recommended Actions
<b>Primary Drive Full</b>			
Primary drive is full.	Major	The primary disk drive reached 100% of usage.	Review the syslog file to identify the possible cause of this problem.
<b>Fan Monitoring</b>			
Fan <n> failure has occurred.	Minor	The specified fan <n> failed, where <n> indicates the fan number.	Inspect the specified fan for mechanical or electrical problems.
Fan <n> is OK.	Informational	The specified fan <n> is functioning properly, where <n> indicates the fan number.	No action is required.
<b>Power Supply Failure: monitored at 1 minute</b>			
A power supply failure has occurred.	Major	The power supply failed.	Inspect the power supply for the possible cause of the failure.
Power supply <n> failure has occurred.	Major	The specified power supply <n> failed, where <n> indicates the power supply number.	Inspect the specified power supply for the possible cause of the failure.
The power supply is OK.	Informational	The power supply is functioning properly.	No action is required.
Power supply <n> is OK.	Informational	The specified power supply <n> is functioning properly, where <n> indicates the power supply number.	No action is required.
<b>RAID monitoring, at 1 minute interval</b>			
A RAID battery failure has occurred.	Major	The system RAID battery failed. The alert light is red.	Inspect the battery for the possible cause of the failure.
The system's RAID battery is OK.	Informational	The system RAID battery is charging and functioning properly. The alert light changed from red to green.	No action is required.
Unable to retrieve RAID array state!	Undetermined	The appliance failed to retrieve the RAID array state. The alert light is red.	Review the syslog file to identify the possible cause of this problem.
The system's RAID array is now running in an optimal state.	Informational	The RAID system is functioning at an optimal state.	No action is required.

The system's RAID array is in a degraded state. RAID Disk <n> is EMPTY.	Major	The RAID system is degrading. The specified RAID Disk <n> is empty, where <n> indicates the RAID disk number.	Review the syslog file to identify the possible cause of this problem.
The system's RAID array is rebuilding. RAID Disk <n> is OFFLINE.	Minor	The RAID system is rebuilding. The specified RAID Disk <n> is offline, where <n> indicates the RAID disk number.	No action is required.
<b>Syslog Backup Processes</b>			
The syslog backup process is successful.	Informational	Rotated syslog files are uploaded successfully to the external backup server.	No action is required.
The syslog backup process failed.	Major	Failed to forward the rotated syslog files to the external backup server.	Review the syslog messages to identify the possible cause of this problem.

## Processing and Software Failure Traps

The `ibSubsystemName` object is associated with certain traps of the Processing and Software Failure traps. Therefore, you cannot map all the traps of the Processing and Software Failure traps with the `ibSubsystemName`. If there is no value in the `ibSubsystemName`, then it belongs to the N/A category. For more information on the values for the `ibSubsystemName`, see the *ibSubsystemName Values* table.

<b>Named Daemon Failure</b>			
A named daemon monitoring failure has occurred.	Critical	The named process failed.	Review the syslog file to identify the possible cause of this problem.
<b>DHCP Daemon Failure</b>			
A DHCP daemon monitoring failure has occurred.	Critical	The dhcpd process failed.	Review the syslog file to identify the possible cause of this problem.
<b>SSH Daemon Failure</b>			
An SSH daemon failure has occurred.	Major	The sshd process failed.	Review the syslog file to identify the possible cause of this problem.
<b>NTP Daemon Failure, monitored every 10 minutes</b>			
An NTP daemon failure has occurred.	Major	The ntpd process failed.	Review the syslog file to identify the possible cause of this problem.
<b>Cluster Daemon Failure</b>			

A grid daemon failure has occurred.	Critical	The clusterd process failed.	Review the syslog file to identify the possible cause of this problem.
<b>LCD Daemon Failure</b>			
An LCD failure has occurred.	Major	The LCD process failed. The alert light is yellow.	<ol style="list-style-type: none"> <li>1. Inspect the LCD panel for the possible cause of this problem.</li> <li>2. Review the syslog file to identify the possible cause of this problem.</li> </ol>
<b>Apache Software httpd failure, monitored every 2 minutes</b>			
An Apache software failure has occurred.	Critical	The request to monitor the Apache server failed.	Review the syslog file to identify the possible cause of this problem.
<b>Serial Console Failure</b>			
An Infoblox serial console software failure has occurred.	Major	The Infoblox serial console failed.	Review the syslog file to identify the possible cause of this problem.
<b>Controld Software Failure</b>			
A controld failure has occurred.	Critical	The controld process failed.	Review the syslog file to identify the possible cause of this problem.
<b>SNMP Sub-agent Failure</b>			
An SNMP server failure has occurred.	Major	The one-subagent process failed.	Review the syslog file to identify the possible cause of this problem.
<b>TFTPD and FTPD Failure</b>			
A TFTP daemon failure has occurred.	Critical	The tftpd process failed.	Review the syslog file to identify the possible cause of this problem.
An FTP daemon failure has occurred.	Critical	The ftpd process failed.	Review the syslog file to identify the possible cause of this problem.
<b>HTTP File Distribution, monitored at 10 second intervals</b>			
An HTTP file distribution daemon failure has occurred.	Critical	The HTTP file distribution process failed.	Review the syslog file to identify the possible cause of this problem.

<b>DNS ONE quagga Processes (zebra &amp; ospfd)</b>			
An OSPF routing daemon failure has occurred.	Critical	Either the zebra process or the ospfd process failed. Both the zebra and ospfd process belongs to ospf subsystem.	Review the syslog file to identify the possible cause of this problem.
<b>Backup Failure</b>			
Backup failed.	Minor	<p>The backup failed. One of the following could be the cause of the failure:</p> <ul style="list-style-type: none"> <li>• The appliance could not access a backup directory.</li> <li>• The backup was interrupted by one of the following signals: SIGINT, SIGHUP, or SIGTERM.</li> <li>• Incorrect login or connection failure in an FTP backup.</li> <li>• The backup failed to create temporary files.</li> </ul>	Review the syslog file to identify the possible cause of this problem.
<b>Database Backup Failure</b>			
Database backup failed.	Not implemented	The db_dump process failed.	Review the syslog file to identify the possible cause of this problem.
<b>Backup Module Failure</b>			
Module backup failed.	Not implemented	The backup of product- specific files failed.	Review the syslog file to identify the possible cause of this problem.
<b>Backup File Size Exceeded</b>			
File size exceeded the quota. Backup failed.	Not implemented	The backup failed because the file size exceeded the limit of 5GB.	Limit the size of the backup file to less than 5GB.
Another backup is in progress. Backup will not be performed.	Not implemented	The backup failed because of an attempt to back up or merge files while another backup or restore was in progress.	Wait until the backup or restore is complete before starting another backup.
<b>Watchdog Process Monitoring</b>			

WATCHDOG: <registered client name> failed on <server IP address>	Critical	The watchdog process detected a registered client failure on a specific server. The <registered client name> could be one of the following: <ul style="list-style-type: none"> <li>• Clusterd_timeout</li> <li>• DB_Sentinel</li> <li>• Process_Manager</li> <li>• Clusterd_monitor</li> <li>• Disk_monitor</li> </ul>	Review the syslog file to identify the possible cause of this problem.
<b>Microsoft Server</b>			
Microsoft server <hostname>/<IP address> has failed.	Major	The Microsoft server could not be reached.	Check that the Microsoft server is connected to the network and configured properly.
Microsoft server <hostname>/<IP address> is OK.	Informational	The Microsoft server can be reached and is functioning properly.	No action is required.
<b>Microsoft DNS/DHCP Service</b>			
Service connection to Microsoft DNS server <hostname>/<IP address> has failed.	Major	The Microsoft DNS service is not responding.	Check that the DNS service is configured and running on the Microsoft server.
Service connection to Microsoft DHCP server <hostname>/<IP address> has failed.	Major	The Microsoft DHCP service is not responding.	Check that the DHCP service is configured and running on the Microsoft server.
Service connection to Microsoft DNS server <hostname>/<IP address> is OK.	Informational	The Microsoft DNS service is responding.	No action is required.
Service connection to Microsoft DHCP server <hostname>/<IP address> is OK.	Informational	The Microsoft DHCP service is responding.	No action is required.
<b>NAC Authentication Server Group</b>			
NAC Authentication server group is down	Major	None of the servers in the NAC authentication server group can be reached.	Review the syslog.
NAC Authentication server group is up	Informational	The NAC authentication server group is responding.	No action is required.
<b>GUI Login</b>			
A GUI login failure has occurred	Major	An admin failed to log in to the GUI.	Check the credentials of the admin.

<b>Serial Console Login</b>			
A Serial Console login failure has occurred	Major	An admin failed to log in through the serial console.	Check the credentials and permissions, and check that the serial console is enabled.
<b>Reboot</b>			
The system is rebooting.	Informational	A system reboot command was sent.	No action is required.
<b>DHCP Lease Conflict</b>			
DHCP address conflicts with an existing lease.	Major	The discovery process found a DHCP lease conflict.	In the IP Map or List panel, select a conflicting address, and then click <b>Resolve Conflict</b> . For more information about resolving DHCP lease conflicts, see <a href="#">Managing Discovered Data</a> .
<b>DHCP Fixed Address Conflict</b>			
DHCP address conflicts with an existing fixed address.	Major	The discovery process found a fixed address conflict.	In the IP Map or List panel, select a conflicting address, and then click <b>Resolve Conflict</b> . For more information about resolving fixed address conflicts, see <a href="#">Managing Discovered Data</a> .
<b>DHCP Range Conflict</b>			
DHCP address conflicts with an existing range.	Major	The discovery process found a conflict with an existing range.	In the IP Map or List panel, select a conflicting address, and then click <b>Resolve Conflict</b> . For more information about resolving DHCP range conflicts, see <a href="#">Managing Discovered Data</a> .
<b>DHCP Host Conflict</b>			
DHCP address conflicts with an existing host address.	Major	The discovery process found a conflict with an existing host address.	In the IP Map or List panel, select a conflicting address, and then click <b>Resolve Conflict</b> . For more information about resolving host conflicts, see <a href="#">Managing Discovered Data</a> .
<b>DNS Health Check Monitor</b>			
DNS Health Check query failed.	Major	DNS Health Check query has failed.	Review the syslog file to identify the possible cause of this problem.
DNS Health Check query has succeeded	Informational	DNS Health Check query has succeeded.	No action required.

<b>Syslog Daemon Failure</b>			
Syslog daemon is not running.	Critical	Syslog process stopped.	Review the syslog file to identify the possible cause of this problem.
<b>BFD Daemon Failure</b>			
An BFD daemon failure has occurred.	Major	BFD has failed to detect failure in the bidirectional path between two interfaces.	Review the syslog file to identify the possible cause of this problem.
<b>Process Stop/Start</b>			
The system stopped and started a process.	Major	The system restarted a process.	Review the syslog file to identify the possible cause of this problem.
<b>Zone Transfer Failed</b>			
A zone transfer failure has occurred.	Major	A zone transfer failed.	Review the syslog file to identify the possible cause of this problem.
<b>RPZ Refresh Failed</b>			
RPZ refresh failure has occurred.	Critical	An RPZ refresh has failed. The appliance sends this trap only when RPZ refresh from all the configured primary servers fail.	Review the syslog file to identify the possible cause of this problem.
<b>Unbound Software Failed</b>			
A unbound daemon monitoring failure has occurred.	Major	Unbound software failure has occurred.	Review the syslog file to identify the possible cause of this problem.
<b>Threat Analytics Service Failed</b>			
Threat Analytics Auto Download has failed.	Major	Auto download for Threat Analytics has failed.	Review the syslog file to identify the possible cause of this problem.
<b>Outbound Service Failed</b>			
The Outbound worker failed.	Major	The Outbound service failure has occurred.	Review the syslog file to identify the possible cause of this problem.
<b>Clear</b>			
SNMP Trap is cleared. LCD failure	N/A	The SNMP Trap for LCD failure is cleared.	No action is required.



SNMP Trap is cleared. Serial Console	N/A	The SNMP Trap for Serial Console is cleared.	No action is required.
SNMP Trap is cleared. ControlD failure	N/A	The SNMP Trap for ControlD failure is cleared.	No action is required.
SNMP Trap is cleared. GUI Login	N/A	The SNMP Trap for GUI Login is cleared.	No action is required.
SNMP Trap is cleared. Serial Console Login	N/A	The SNMP Trap for Serial Console Login is cleared.	No action is required.
SNMP Trap is cleared. SSHD failure	N/A	The SNMP Trap for SSHD failure is cleared.	No action is required.
SNMP Trap is cleared. LDAP servers	N/A	The SNMP Trap for LDAP server is cleared.	No action is required.
SNMP Trap is cleared. OCSP Responders	N/A	The SNMP Trap for OCSP Responders is cleared.	No action is required.
SNMP Trap is cleared. OSPF	N/A	The SNMP Trap for OSPF is cleared.	No action is required.
SNMP Trap is cleared. OSPF6	N/A	The SNMP Trap for OSPF6 is cleared.	No action is required.
SNMP Trap is cleared. BGP	N/A	The SNMP Trap for BGP is cleared.	No action is required.
SNMP Trap is cleared. HSM	N/A	The SNMP Trap for HSM is cleared.	No action is required.
SNMP Trap is cleared. HTTP	N/A	The SNMP Trap for HTTP is cleared.	No action is required.
SNMP Trap is cleared. Cluster	N/A	The SNMP Trap for Cluster is cleared.	No action is required.
SNMP Trap is cleared. DuplicateIP	N/A	The SNMP Trap for Duplicate IP is cleared.	No action is required.
<b>Restart</b>			
The system is being restarted.	Informational	A system restart command was sent.	No action is required.
<b>DNS Integrity Check</b>			

<p>Cannot perform DNS Integrity Check because the appliance is unable to connect to the external DNS server. There are list of nameservers failure: &lt;IP addresses&gt;</p>	<p>N/A</p>	<p>The DNS integrity check cannot be performed because the appliance is unable to connect to the external DNS server.</p>	<p>No action is required.</p>
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------	---------------------------------------------------------------------------------------------------------------------------	-------------------------------

**Threshold Crossing Traps**

<b>ibTrapDesc</b> <b>OID 3.1.1.1.2.11.0</b>	<b>ibTrapSeverity</b>	<b>ibObjectName</b> <b>OID 3.1.1.1.2.3.0</b>	<b>Description/Cause</b>	<b>Recommended Actions</b>	
<p><b>System Memory Usage</b></p>					

System has run out of memory.	Major	memory	<p>The appliance ran out of memory. The appliance encountered this problem when one of the following occurred:</p> <ul style="list-style-type: none"> <li>• The total free memory on the appliance was less than or equal to 0%.</li> <li>• The total physical memory was less than the total free memory.</li> <li>• The percentage of free memory compared to the total physical memory was less than 5%, and the free swap percentage was less than 80%.</li> <li>• The percentage of free memory compared to the total physical memory was less than 5%, plus the numbers of both swap INs and swap OUTs were greater than or equal to 3,200.</li> </ul>		
-------------------------------	-------	--------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--	--

- The percentage of free memory compared to the total physical memory was between 5% and 10%, the free swap percentage was greater than or equal to 80%, plus the numbers of both swap INs and swap OUTs were greater than or equal to 3,200.
- The percentage of free memory compared to the total physical memory was greater than 10%, the free swap percentage was less than 80%, plus the numbers of both swap INs and swap OUTs were greater than or equal to 3,200.

**Note:** Free memory = free physical RAM + free cache buffers. The high threshold for swap pages is 3,200.

<p>System memory usage exceeds the critical threshold value.</p>	<p>Minor</p>	<p>memory</p>	<p>The memory usage on the appliance exceeded the configured Trigger value. For more information about defining thresholds for traps, see <a href="#">Configuring SNMP</a>. The appliance encountered this problem when one of the following occurred:</p> <ul style="list-style-type: none"> <li>• The percentage of free memory compared to the total physical memory was less than 5%, and the free swap percentage was less than 90%.</li> <li>• The percentage of free memory compared to the total physical memory was less than 5%, plus the number of swap INs was less than 3,200 and the number of swap OUTs was greater than or equal to 3,200.</li> </ul>		
------------------------------------------------------------------	--------------	---------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--	--

			<ul style="list-style-type: none"> <li>• The percentage of free memory compared to the total physical memory was between 5% and 10%, and the free swap percentage was less than 80%.</li> <li>• The percentage of free memory compared to the total physical memory was greater than 5%, plus the number of swap INs was less than 3,200 and the number of swap OUTs was greater than or equal to 3,200.</li> </ul> <p><b>Note:</b> Free memory = free physical RAM + free cache buffers. The high threshold for swap pages is 3,200.</p>	
System memory usage is OK.	Minor	memory	The memory usage on the system is at or below the Reset value after it went above the Trigger value.	
<b>Primary Hard Drive Usage (monitored every 30 seconds)</b>				

System primary hard disk usage is over threshold value.	Minor	disk_usage	The appliance sends this trap when primary hard disk usage first exceeds the configured Trigger value. The default value is 85. The alert light is yellow.  For more information about defining thresholds for traps, see <a href="#">Configuring SNMP</a> .	Review the syslog file to identify the possible cause of this problem.
Primary drive is full.	Major	disk_usage	The primary hard disk usage exceeded 95%. The alert light is red.	Review the syslog file to identify the possible cause of this problem.
Primary drive usage is OK.	Minor	disk_usage	The appliance sends this trap when the primary hard disk usage first moves at or below the configured Reset value after it exceeded the Trigger value. The default is 70. The alert light is green.	No action is required.
<b>CPU Usage</b>				
CPU usage above threshold value.	Major	cpu_usage	CPU usage exceeded the Trigger value for 15 seconds. For more information about defining thresholds for traps, see <a href="#">Configuring SNMP</a> .	Monitor CPU usage.
CPU usage OK.	Minor	cpu_usage	CPU usage dipped below the reset value after the "CPU usage above threshold value" trap was sent.	No action is required.
<b>Note:</b> Use the CLI command <code>set thresholdtrap</code> to enable the CPU usage trap and configure the trigger and reset values. For information, refer to the <i>Infoblox CLI Guide</i> .				
<b>Swap Usage</b>				
System swap space usage exceeds the critical threshold value.	Major	swap_usage	System swap space usage exceeded the Trigger value. For more information about defining thresholds for traps, see <a href="#">Configuring SNMP</a> .	Monitor System swap usage.



System swap space usage is OK.	Minor	swap_usage	System swap space usage dipped below the reset value after it exceeded the Trigger value.	No action is required.
<b>Replication Statistics Monitoring</b>				
Grid queue replication problem.	N/A	For send trap: Cluster_Send_Queue For receive trap: Cluster_Recv_Queue	The system encountered this problem when all of the following conditions occurred: <ul style="list-style-type: none"> <li>• The node was online.</li> <li>• The number of the replication queue being sent from the master column was greater than 0, or the number of the queue received was greater than 0.</li> <li>• It was more than 10 minutes since the last replication queue was sent and monitored.</li> </ul>	Review the syslog file to identify the possible cause of this problem.
<b>DHCP Range Threshold Crossing</b>				

<p>DHCP high threshold crossed:  Member:  &lt;DHCP server node VIP&gt;  Network:  &lt;network&gt;/  &lt;network view&gt;  Range: &lt;DHCP range&gt;/ &lt;network view&gt;  High Trigger Mark:  &lt;high percentage&gt;  (95% by default) High  Reset Mark:  &lt;reset  percentage&gt; (80% by  default)  Current Usage:  &lt;current usage  percentage&gt;  Active Leases:  &lt;number of active  leases&gt;  Available Leases:  &lt;number of available  leases&gt;  Total Addresses:  &lt;total addresses&gt;</p>	N/A	Threshold	The system encountered this problem when the address usage in the DHCP range is greater than the configured High Trigger value.	Review the syslog file to identify the possible cause of this problem.
<p>DHCP high threshold reset:  Member:  &lt;DHCP server node VIP&gt;  Network:  &lt;network&gt;/  &lt;network view&gt;  Range: &lt;DHCP range&gt;/&lt;network view&gt;  High Trigger Mark:  &lt;high percentage&gt;  (95% by default) High  Reset Mark:  &lt;reset  percentage&gt; (80% by  default)  Current Usage:  &lt;current usage  percentage&gt;  Active Leases:  &lt;number of active  leases&gt;  Available Leases:  &lt;number of available  leases&gt;  Total Addresses:  &lt;total addresses&gt;</p>	N/A	Threshold	The system encountered this problem when the address usage in the DHCP range goes below the Reset value after it hit the Trigger value.	
<b>IPAM Utilization Threshold Crossing</b>				
<p>Network IPAM Utilization capacity usage is over the threshold value.  Network:  &lt;network&gt;/  &lt;network view&gt;</p>	N/A	Threshold	The appliance sends this trap when the IPAM utilization for a network is above the configured Trigger value.	Review the syslog file to identify the possible cause of this problem.

Network Container IPAM Utilization capacity usage is over the threshold value. Network container: <network>/ <network view>	N/A	Threshold	The appliance sends this trap when the IPAM utilization for a network container is above the configured Trigger value.	Review the syslog file to identify the possible cause of this problem.
Network IPAM Utilization capacity usage is OK. Network: <network>/ <network view>	N/A	Threshold	The appliance sends this trap when the IPAM utilization for a network is below the configured Reset value.	Review the syslog file to identify the possible cause of this problem.
Network Container IPAM Utilization capacity usage is OK. Network container: <network>/ <network view>	N/A	Threshold	The appliance sends this trap when the IPAM utilization for a network container is below the configured Reset value.	Review the syslog file to identify the possible cause of this problem.
<b>DHCP DDNS Updates Deferred</b>				
DHCP DNS updates deferred: Retried at least once: <number of retries> Maximum number of deferred updates since start of problem episode (or restart): <max number>	N/A	Threshold	The DNS updates were deferred because of DDNS update errors.	Review the syslog file to identify the possible cause of this problem.
<b>RPZ Hit Rate</b>				
RPZ hit rate exceed normal value.	N/A	Threshold	The appliance sends this trap when the RPZ hit rate exceeds the configured Trigger value.	Review the syslog file to identify the possible cause of this problem.
RPZ hit rate has returned to normal value.	N/A	Threshold	The appliance sends this trap when the RPZ hit rate equals the configured Reset value.	Review the syslog file to identify the possible cause of this problem.
<b>Threat Protection Total Traffic</b>				
Threat Protection Service total traffic is above threshold.	N/A	Threshold	The appliance sends this trap when the Treat Protection total traffic exceeds the configured Trigger value.	Review the syslog file to identify the possible cause of this problem.

Threat Protection Service total traffic is OK.	N/A	Threshold	The appliance sends this trap when the Threat Protection total traffic is less than the configured Reset value.	Review the syslog file to identify the possible cause of this problem.
<b>Threat Protection Dropped Traffic</b>				
Threat Protection Service dropped traffic is above threshold.	N/A	Threshold	The appliance sends this trap when the Treat Protection dropped traffic exceeds the configured Trigger value.	Review the syslog file to identify the possible cause of this problem.
Threat Protection Service dropped traffic is OK.	N/A	Threshold	The appliance sends this trap when the Threat Protection dropped traffic is less than the configured Reset value.	Review the syslog file to identify the possible cause of this problem.
<b>Database Capacity Usage</b>				
Database capacity used is over the threshold value.	Minor	db_usage	The appliance database usage exceeded the configured threshold value.	Increase the database capacity.
Database capacity used is OK.	Minor	db_usage	The appliance database usage is less than the configured threshold value.	No action is required.
<b>DNS Monitor</b>				

<p>DNS security alert. There were <i>actual</i> DNS responses to invalid ports in the last minute, comprising <i>percent%</i> of all responses. Primary sources: <i>ip_address</i> sent <i>count</i>, <i>ip_address</i> sent <i>count</i>.</p>	<p>Major</p>	<p>dns_security_port</p>	<p>DNS security alert. There were <i>actual</i> DNS responses to invalid ports in the last minute, comprising <i>percent%</i> of all responses. Primary sources: <i>ip_address</i> sent <i>count</i>, <i>ip_address</i> sent <i>count</i>. where</p> <ul style="list-style-type: none"> <li>• <i>actual</i> is the total number of DNS responses arrive on invalid ports.</li> <li>• <i>percent%</i> is the percentage of invalid DNS responses over the total number of DNS responses.</li> <li>• <i>ip_address</i> is the IP address of the primary source that generated the invalid DNS responses.</li> <li>• <i>count</i> is the number of invalid responses generated by the specified IP address.</li> </ul>	<ol style="list-style-type: none"> <li>1. Review the following: <ul style="list-style-type: none"> <li>• DNS alert status</li> <li>• syslog file</li> </ul> </li> <li>2. Limit access or block connections from the primary sources. For information about configuring rate limiting rules, see <a href="#">Monitoring Tools</a>.</li> </ol>
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--------------	--------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

			<p>Example: DNS security alert. There were 1072 DNS responses to invalid ports in the last minute, comprising 92% of all responses. Primary sources: 10.0.0.0 sent 1058, 2.2.2.2 sent 14.</p>	
--	--	--	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--

<p>DNS security alert. There were <i>actual</i> DNS responses with invalid TXID in the last minute, comprising <i>percent%</i> of all responses. Primary sources: <i>ip_address</i> sent <i>count</i>, <i>ip_address</i> sent <i>count</i>.</p>	<p>N/A</p>	<p>dns_security_t_xid</p>	<p>DNS security alert. There were <i>actual</i> DNS responses with invalid TXID in the last minute, comprising <i>percent%</i> of all responses. Primary sources: <i>ip_address</i> sent <i>count</i>, <i>ip_address</i> sent <i>count</i>. where</p> <ul style="list-style-type: none"> <li>• <i>actual</i> is the total number of DNS responses that have invalid TXIDs.</li> <li>• <i>percent%</i> is the percentage of invalid DNS responses over the total number of DNS responses.</li> <li>• <i>ip_address</i> is the IP address of the primary source that generated the invalid DNS responses.</li> <li>• <i>count</i> is the number of invalid responses generated by the specified IP address.</li> </ul>	<ol style="list-style-type: none"> <li>1. Review the following: <ul style="list-style-type: none"> <li>• DNS alert status</li> <li>• syslog file</li> </ul> </li> <li>2. Limit access or block connections from the primary sources. For information about configuring rate limiting rules, see <a href="#">Monitoring Tools</a>.</li> </ol>
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------	---------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

			Example: DNS security alert. There were 1072 DNS responses with invalid TXID in the last minute, comprising 92% of all responses. Primary sources: 10.0.0.0 sent 1058, 2.2.2.2 sent 14.		
<b>RootFS Partition Monitor</b>					
Root file system is full.	Major	Root filesystem	The Root filesystem usage exceeded the maximum.	Review the syslog file to identify the possible cause of this problem.	
Root file system disk usage is over threshold value.	Minor	Root filesystem	The appliance sends this trap when the Root filesystem usage first exceeds the configured Trigger value.	Review the syslog file to identify the possible cause of this problem.	
Root file system disk usage is OK.	Minor	Root filesystem	The appliance sends this trap when the Root filesystem disk usage first moves at or below the configured Reset value after it exceeded the Trigger value. For more information about defining thresholds for traps, see <a href="#">Configuring SNMP</a> .	No action	
<b>Reporting</b>					
Reporting drive is full.	Major	reporting_service	Reporting drive reached the maximum capacity.	Review the syslog file to identify the possible cause of this problem.	



Reporting drive usage is over threshold value.	Minor	reporting_service	The appliance sends this trap when the Reporting volume first exceeds the configured Trigger value. The default Trigger value is 80.	Review the syslog file to identify the possible cause of this problem.	
Reporting drive usage is OK.	Minor	reporting_service	The appliance sends this trap when the Reporting volume first moves at or below the configured Reset value after it exceeded the Trigger value. The default Reset value is 71. For information on setting the Trigger and Reset values, see <a href="#">Configuring SNMP</a> .	No action	
<b>File Distribution</b>					
File Distribution services storage usage reached the threshold value.	N/A	File Distribution	File distribution service storage reached the configured threshold value.	Review the syslog file to identify the possible cause of this problem.	

## Object State Change Traps

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	Description/Cause	Recommended Actions
<b>Service Shutdown</b>			
Shutting down services due to database snapshot.	Major	The appliance is shutting down its services while synchronizing the database with the Grid Master.	No action is required.
<b>Network Interfaces Monitoring</b>			
LAN1 port link is down. Please check the connection.	Major	The LAN1 port is up, but the link is down.	Check the LAN1 link connection.
LAN2 port link is down. Please check the connection.	Major	The LAN2 port is up, but the link is down.	Check the LAN2 link connection.
HA port link is down. Please check the connection.	Major	The HA port is up, but the link is down.	Check the HA link connection.
MGMT port link is down. Please check the connection.	Major	The MGMT port is enabled, but the link is down.	Check the MGMT link connection.
LAN1 port link is up.	Major	The LAN1 port link is up and running.	No action is required.

LAN2 port link is up.	Major	The LAN2 port link is up and running.	No action is required.
HA port link is up.	Major	The HA port link is up and running.	No action is required.
MGMT port link is up.	Major	The MGMT port link is up and running.	No action is required.
<b>HA State Change from Initial to Active</b>			
The node has become ACTIVE.	Informational	A node in an HA pair becomes active. The HA pair starts up.	No action is required.
<b>HA State Change from Passive to Active</b>			
The node has become ACTIVE.	Informational	The node changed from a passive to an active node.	No action is required.
<b>HA State Change to Active-Active</b>			
The node is in an ACTIVE-ACTIVE state.	Informational	The node is in the active state.	No action is required.
<b>HA State Change from Initial to Passive</b>			
The node has become PASSIVE.	Informational	A node in an HA pair becomes passive. The HA pair starts up, and the node is not a Grid Master candidate.	No action is required.
<b>Node Connected to Grid</b>			
The Grid member is connected to the Grid master.	Informational	The Grid member joined the Grid, and it is not a Grid Master candidate.	No action is required.
<b>Node Disconnected from Grid</b>			
The Grid member is not connected to the Grid master.	Informational	The Grid member lost its connection to the Grid Master.	No action is required.
<b>Replication State Monitoring</b>			
HA replication is online.	Informational	The HA replication is online.	No action is required.
HA replication is offline.	Informational	The HA replication is offline.	No action is required.
<b>NTP is out of sync, monitored every 30 seconds</b>			
The NTP service is out of synchronization.	Major	The Infoblox NTP server and the external NTP server are not synchronized.	Review the syslog file to identify the possible cause of this problem.

NTP Service is working.	Informational	The NTP service started working again.	No action is required.
<b>NTP service is synchronizing to local clock</b>			
The NTP service is synchronizing to local clock.	Major	The Infoblox NTP server synchronizes its clients with its local clock.	No action is required.
<b>DHCP service state change</b>			
DHCP Service is working.	Informational	The DHCP service started working again.	No action is required.
DHCP Service is in a warning state.	Informational	The DHCP service is in a warning state.	Review the syslog file
DHCP Service Failed	Informational	The DHCP service is in failed state.	Review the syslog file.
DHCP Service is inactive.	Informational	The DHCP service became inactive.	Check if an admin disabled the service.
<b>DNS service state change</b>			
DNS Service is working.	Informational	The DNS service started working again.	No action is required.
DNS Service failed	Informational	The DNS service is in a failed state.	Review the syslog file.
DNS Service is inactive.	Informational	The DNS service became inactive.	Check if an admin disabled the service.
<b>NTP service state change</b>			
The NTP service resumed synchronization.	Informational	The NTP service started working again.	No action is required.
NTP Service is inactive.	Informational	The NTP service became inactive.	Check if an admin disabled the service.
<b>TFTP service state change</b>			
TFTP Service is working.	Informational	The TFTP service started working again.	No action is required.
TFTP Service is inactive.	Informational	The TFTP service became inactive.	Check if an admin disabled the service.
<b>FTP service state change</b>			
FTP Service is working.	Informational	The FTP service started working again.	No action is required.
FTP Service is inactive.	Informational	The FTP service became inactive.	Check if an admin disabled the service.

<b>HTTP service state change</b>			
HTTP File Dist Service is working.	Informational	The HTTP file distribution service started working again.	No action is required.
HTTP File Dist Service is inactive.	Informational	The HTTP file distribution service became inactive.	Check if an admin disabled the service.
<b>bloxTools service state change</b>			
BloxTools Service is working.	Informational	The bloxTools service started working again.	No action is required.
BloxTools Service is in warning state.	Informational	The bloxTools service is in a warning state.	Review the syslog file
BloxTools Service is inactive.	Informational	The bloxTools service became inactive.	Check if an admin disabled the service.
BloxTools Service failed.	Critical	The bloxTools daemon failed.	Review the syslog file.
<b>Captive Portal service state change</b>			
Captive Portal Service is working.	Informational	The captive portal service started working again.	No action is required.
Captive Portal Service is inactive.	Informational	The captive portal service became inactive.	Check if an admin disabled the service.
<b>Discovery Collector service state change</b>			
Discovery Collector Service is working.	Informational	The discovery collector service started working again.	No action is required.
Discovery Collector Service is inactive.	Informational	The discovery collector service became inactive.	Check if an admin disabled the service.
Discovery Collector Service is in warning state.	Informational	The discovery collector service is in a warning state.	Review the syslog file.
Discovery Collector Service has failed.	Informational	The discovery collector service has failed.	Review the syslog file.
<b>Discovery Collector service state change</b>			
Discovery Consolidator Service is working.	Informational	The discovery consolidator service started working again.	No action is required.
Discovery Consolidator Service is inactive.	Informational	The discovery consolidator service became inactive.	Check if an admin disabled the service.
Discovery Consolidator Service is in warning state.	Informational	The discovery consolidator service is in a warning state.	Review the syslog file

Discovery Consolidator Service has failed.	Critical	The discovery consolidator service has failed.	Review the syslog file
<b>IF-MAP service state change</b>			
IFMAP Service is inactive.	Informational	The IF-MAP service became inactive.	Check if an admin disabled the service.
IFMAP Service is working.	Informational	The IF-MAP service started working again.	No action is required.
<b>LDAP service state change</b>			
LDAP service is inactive.	Informational	The LDAP service became inactive.	Check if an admin disabled the LDAP service.
All LDAP servers are available.	Informational	All LDAP servers are available.	No action is required.
At least one LDAP server is unavailable.	Major	At least one LDAP server is out of service.	<ul style="list-style-type: none"> <li>• Check if an admin disabled the LDAP server.</li> <li>• Review the syslog file.</li> </ul>
All LDAP servers are unavailable.	Informational	All LDAP servers are not available.	<ul style="list-style-type: none"> <li>• Check if an admin disabled the LDAP server.</li> <li>• Review the syslog file.</li> </ul>
<b>Certificate Authentication Service state change</b>			
CAS service is working.	Informational	The certificate authentication service started working again.	No action is required.
CAS service is inactive.	Informational	The certificate authentication service became inactive.	Check if an admin disabled the service.
<b>OCSP responders state change</b>			
All OCSP responders available.	Informational	The OCSP responders are available.	No action is required.
All OCSP responders are out of service.	Informational	The OCSP responders are out of service.	<ul style="list-style-type: none"> <li>• Check if an admin disabled the service.</li> <li>• Review the syslog file.</li> </ul>

OCSP service inactive.	Informational	The OCSP responders service became inactive.	<ul style="list-style-type: none"> <li>• Check if an admin disabled the service.</li> <li>• Review the syslog file.</li> </ul>
<b>Reporting service state change</b>			
Reporting Service is working.	Informational	The Reporting service started working again.	No action is required.
Reporting Service failed.	Informational	The Reporting Service is in Failed state.	Review the syslog file.
Reporting Service is inactive.	Informational	The Reporting service became inactive.	Check if an admin disabled the service.
Reporting Service is in warning state.	Informational	The Reporting service is in warning state.	Review the audit log file.
<b>Threat Analytics state change</b>			
Threat Analytics Service is working.	Informational	The Threat Analytics service started working again.	No action is required.
Threat Analytics Service is inactive.	Informational	The Threat Analytics service became inactive.	Check if an admin disabled the service.
Threat Analytics Service is failed.	Informational	The Threat Analytics service has failed.	Review the audit log file.
Threat Analytics Service is in warning state.	Informational	The Threat Analytics service is in warning state.	Review the audit log file.
<b>RPZ refresh state change</b>			
RPZ refresh is OK.	Informational	An RPZ refresh succeeded. The appliance sends this trap every time an RPZ zone transfer is successful.	No action is required.
<b>Cisco ISE service state change</b>			
Cisco ISE server is OK.	Informational	The Cisco ISE service started working again.	No action is required.
Connection Error has occurred In Cisco ISE server.	Informational	There is a loss of connection between the NIOS appliance and Cisco ISE server.	<ul style="list-style-type: none"> <li>• Check if an admin disabled the service.</li> <li>• Review the audit log file.</li> </ul>
<b>Cloud API service state change</b>			
Cloud API service is working.	Informational	The Cloud API service started working again.	No action is required.

Cloud API service is inactive.	Informational	The Cloud API service became inactive	<ul style="list-style-type: none"> <li>• Check if an admin disabled the Cloud API service.</li> <li>• Review the audit log file.</li> </ul>
Cloud API service has failed.	Critical	The Cloud API service has failed.	<ul style="list-style-type: none"> <li>• Check if an admin disabled the Cloud API service.</li> <li>• Review the audit log file.</li> </ul>
<b>DNS Integrity Check state change</b>			
DNS Integrity Check is working.	Informational	The DNS Integrity Check started working again.	No action is required.
DNS Integrity Check failed.	Informational	The DNS Integrity Check has failed.	<ul style="list-style-type: none"> <li>• Check if an admin disabled the service.</li> <li>• Review the audit log file.</li> </ul>
DNS Integrity Check severity has changed.	Informational	The DNS Integrity Check has changed.	Review the syslog file.
<b>TAXII service state change</b>			
TAXII service is working.	Informational	The TAXII service started working again.	No action is required.
TAXII Service is failed.	Informational	The TAXII service has failed.	<ul style="list-style-type: none"> <li>• Check if an admin disabled the service.</li> <li>• Review the audit log file.</li> </ul>
TAXII service is inactive.	Informational	The TAXII service became inactive.	<ul style="list-style-type: none"> <li>• Check if an admin disabled the service.</li> <li>• Review the audit log file.</li> </ul>
TAXII Service is in warning state.	Informational	The TAXII service is in warning state.	Review the audit log file.
<b>Outbound service state change</b>			

The Outbound Service Manager stopped.	Informational	The Outbound service manager has stopped working.	<ul style="list-style-type: none"> <li>• Check if an admin disabled the service.</li> <li>• Review the audit log file.</li> </ul>
The Outbound Service Manager failed.	Major	The Outbound service manager has failed.	<ul style="list-style-type: none"> <li>• Check if an admin disabled the service.</li> <li>• Review the audit log file.</li> </ul>
The Outbound Service Manager started.	Informational	The Outbound service manager has started.	No action is required.
The Outbound worker failed.	Major	The Outbound worker has failed.	Review the audit log file.
<b>IPMI Device state change</b>			
IPMI is used by some hardware monitors to test hardware health. The IPMI Device has not responded. Hardware monitors may show spurious failures.	Informational	The IPMI Device has not responded.	<ul style="list-style-type: none"> <li>• Check if an admin disabled the IPMI interface.</li> <li>• Review the audit log file.</li> </ul>
IPMI is used by some hardware monitors to test hardware health. Because the IPMI Device has not responded for a while, hardware monitor failures are likely to be spurious.	Informational	The IPMI Device has not responded.	<ul style="list-style-type: none"> <li>• Check if an admin disabled the IPMI interface.</li> <li>• Review the audit log file.</li> </ul>
IPMI is used by some hardware monitors to test hardware health. The IPMI Device is now available; subsequent hardware monitor failures are likely to be genuine.	Informational	The IPMI Device is now available	No action required.

### Process Started and Stopped Traps

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	Description/Cause	Recommended Actions
<b>Httpd Start</b>			
The process started normally.	Informational	The httpd process started.	No action is required.
<b>Httpd Stop</b>			
The process stopped normally.	Informational	The httpd process stopped.	No action is required.



<b>Process Stop/Start</b>			
The system stopped and started a process.	Major	The system restarted a process.	No action is required.
<b>Zone Transfer Failed</b>			
A zone transfer failure has occurred.	Critical	A zone transfer failed.	Review the syslog file

### Revoked License Trap

ibTrapDesc OID 3.1.1.1.2.11.0	ibTrapSeverity	Description/Cause	Recommended Actions
<b>Revoked License</b>			
This trap is generated when a license is revoked	Critical	A license was revoked.	Obtain and install new license

### ibPlatformOne MIB

The ibPlatformOne MIB provides information about the CPU temperature of the appliance, the replication status, the average latency of DNS requests, DNS security alerts, CPU and memory utilization of the appliance, and the Infoblox service status. The figure below illustrates the structure of the PlatformOne MIB. (Note that the OIDs in the illustration do not include the prefix .1.3.6.1.4.1.7779.)

The ibPlatformOne MIB contains the following objects:

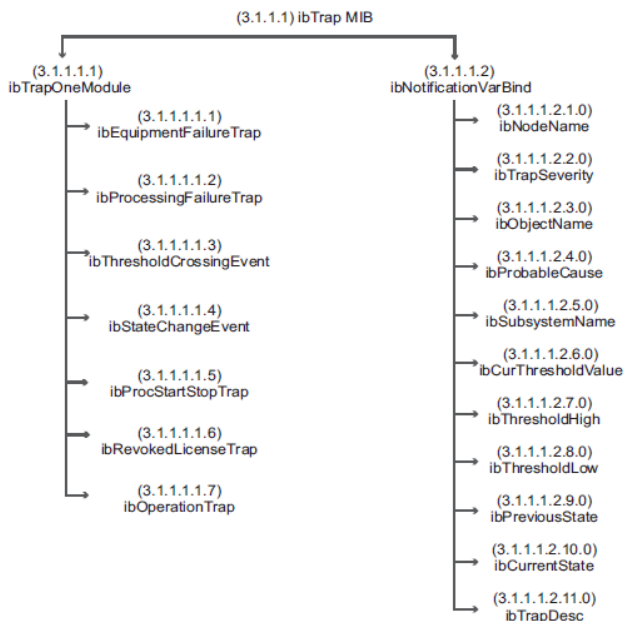
- **ibCPUTemperature (IbString)** tracks the CPU temperature of the appliance.
- **ibClusterReplicationStatusTable** provides information in tabular format about the replication status of the appliance.
- **ibNetworkMonitor** provides information about the average latency of authoritative and nonauthoritative replies to DNS queries for different time intervals. It also provides information about invalid DNS responses that arrive on invalid ports or have invalid DNS transaction IDs.
- **ibHardwareType (IbString)** provides information about the hardware platform. For an Infoblox appliance, it provides the model number of the Infoblox hardware platform. For vNIOS appliances, it identifies if the hardware platform is VMware.
- **ibHardwareId (IbString)** provides the hardware ID of the NIOS appliance.
- **ibSerialNumber (IbString)** provides the serial number of the Infoblox hardware platform.
- **ibNiosVersion (IbString)** provides the version of the NIOS software.
- **ibSystemMonitor** provides information about the CPU and memory utilization of the appliance.
- **ibGridStatus** provides information about an appliance. It indicates whether the appliance is a Grid Master, member, or an independent appliance.
- **ibHAStatus** provides information about the HA status of a member. It indicates if the member is part of an HA configuration, and if it is the active or passive node.
- **ibGridMasterCandStatus** indicates if a member is a Grid Master candidate.
- **ibGridMasterVIP** provides the Grid Master virtual IP address.
- **ibGridReplicationState** provides information about the replication status.

The ibPlatformOne MIB also contains the following tables that provide status of the Infoblox services as well as system and hardware services on the appliance you query:

- **ibMemberServiceStatusTable** provides status of the Infoblox services, such as the DNS and DHCP services, on a queried appliance.
- **ibMemberNodeServiceStatusTable** provides status of the system and hardware services on a queried appliance.

- `ibMemberPassiveNodeServiceStatusTable` provides status of the system and hardware services on the passive node of an HA pair if the queried appliance is the VIP or the active node of an HA pair. For independent appliances and the passive nodes of HA pairs, this table does not display any status.

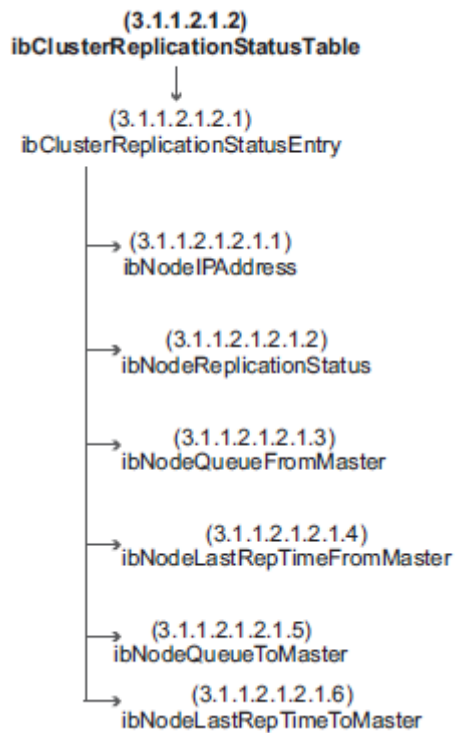
### *ibPlatformOne MIB Structure*



### `ibClusterReplicationStatusTable`

`ibClusterRepliacionStatusTable` (object ID 3.1.1.2.1.2.1) provides information about the Grid replication status. The following figure shows the sub branches of `ibClusterReplicationStatusTable`:

### *ibClusterReplicationStatusTable Objects*



The following table provides information about the `ibClusterReplicationStatusTable` objects:

*ibClusterReplicationStatusTable* Objects

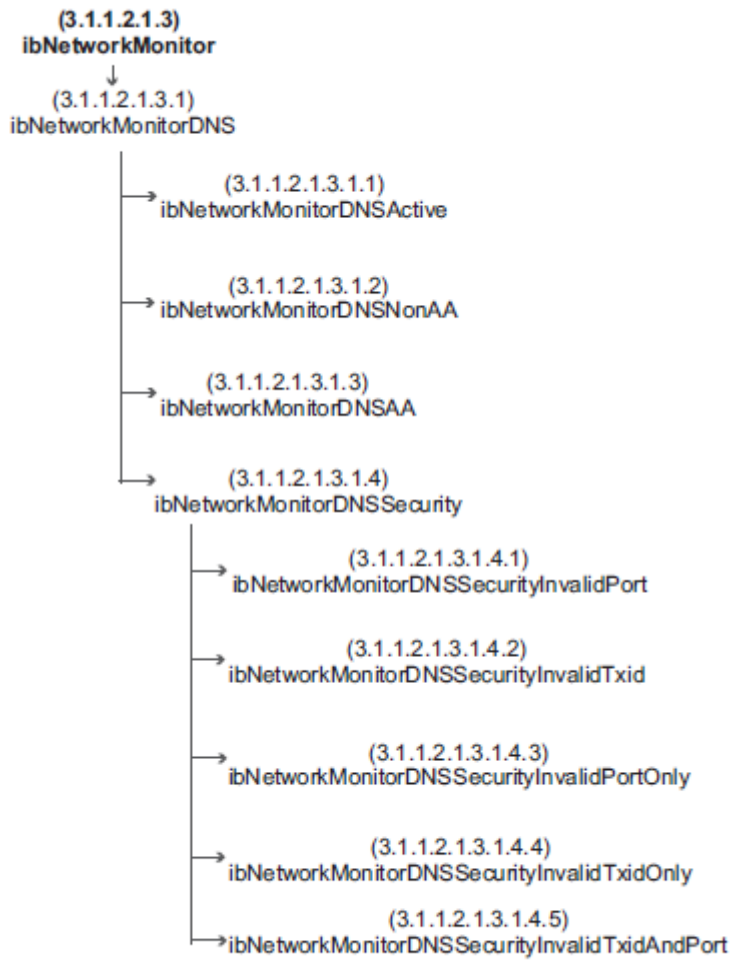
Object (Type)	Description
<code>ibClusterReplicationStatusEntry</code>	A conceptual row that provides information about the Grid replication status. The status indicates whether the appliance is sending replication queues, receiving queues, or having problems with the replication.
<code>ibNodeIPAddress (IbIpAddr)</code>	IP address of a Grid member.
<code>ibNodeReplicationStatus (IbString)</code>	Replication status of the Grid member. The replication status can be one of the following: online, offline, or snapshotting.
<code>ibNodeQueueFromMaster (Integer)</code>	"Sent" queue size from master.
<code>ibNodeLastRepTimeFromMaster (IbString)</code>	Last sent time from master.
<code>ibNodeQueueToMaster (Integer)</code>	"Receive" queue size from master.
<code>ibNodeLastRepTimeToMaster (IbString)</code>	Last receive time from master.

## ibNetwork Monitor

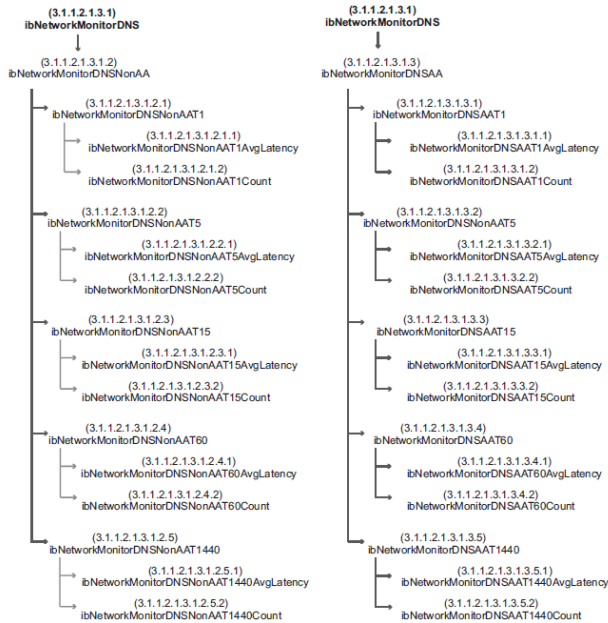
As shown in the following `ibNetWorkMonitorObjects` figure, the `ibNetwork Monitor` has one subtree, `ibNetworkMonitorDNS`, that branches out into the following:

- `ibNetworkMonitorDNSActive` (Integer) reports on whether DNS latency monitoring is enabled. You can enable DNS latency monitoring using the CLI command `set monitor dns`. For more information about enabling and disabling DNS alert monitoring, see [Monitoring Tools](#). This is the only object in this branch. When you send a query for this object, the appliance responds with either "active" (1) or "nonactive" (0).
- `ibNetworkMonitorDNSNonAA` provides information about the average latency of nonauthoritative replies to DNS queries for 1-, 5-, 15-, and 60-minute intervals.
- `ibNetworkMonitorDNSAA` provides information about the average latency of authoritative replies to DNS queries for 1-, 5-, 15-, and 60-minute intervals.
- `ibNetworkMonitorDNSSecurity` provides information about the invalid DNS responses that arrive on invalid ports or have invalid DNS transaction IDs. `ibNetworkMonitorDNSSecurity` branches out into the following:
  - `ibNetworkMonitorDNSSecurityInvalidPort`
  - `ibNetworkMonitorDNSSecurityInvalidTxid`
  - `ibNetworkMonitorDNSSecurityInvalidPortOnly` (Counter)
  - `ibNetworkMonitorDNSSecurityInvalidPortCount` (Counter)
  - `ibNetworkMonitorDNSSecurityInvalidTxidOnly` (Counter)
  - `ibNetworkMonitorDNSSecurityInvalidTxidCount` (Counter)
  - `ibNetworkMonitorDNSSecurityInvalidTxidAndPort` (Counter)

### *ibNetWorkMonitorObjects*



*ibNetworkMonitorDNSNonAAandibNetworkMonitorDNSAASubtrees*



The table below describes the objects in `ibNetworkMonitorDNSNonAA`. You can send queries to retrieve values for these objects.

### *ibNetworkMonitorDNSNonAA Objects*

Object (Type)	Description
<code>ibNetworkMonitorDNSNonAAAT1</code>	File that contains the objects for monitoring the average latency of nonauthoritative replies to queries in the last minute.
<code>ibNetworkMonitorDNSNonAAAT1AvgLatency</code> (Integer)	Indicates the average latency in microseconds of nonauthoritative replies to queries in the last minute.
<code>ibNetworkMonitorDNSNonAAAT1Count</code> (Integer)	Indicates the number of queries used to calculate the average latency of nonauthoritative replies in the last minute.
<code>ibNetworkMonitorDNSNonAAAT5</code>	File that contains the objects for monitoring the average latency of nonauthoritative replies to queries in the last five minutes.
<code>ibNetworkMonitorDNSNonAAAT5AvgLatency</code> (Integer)	Indicates the average latency in microseconds of nonauthoritative replies to queries in the last five minutes.
<code>ibNetworkMonitorDNSNonAAAT5Count</code> (Integer)	Indicates the number of queries used to calculate the average latency of nonauthoritative replies in the last five minutes.
<code>ibNetworkMonitorDNSNonAAAT15</code>	File that contains the objects for monitoring the average latency of nonauthoritative replies to queries in the last 15 minutes.
<code>ibNetworkMonitorDNSNonAAAT15AvgLatency</code> (Integer)	Indicates the average latency in microseconds of nonauthoritative replies to queries in the last 15 minutes.

ibNetworkMonitorDNSNonAAT15Count (Integer)	Indicates the number of queries used to calculate the average latency of nonauthoritative replies in the last 15 minutes.
ibNetworkMonitorDNSNonAAT60	File that contains the objects for monitoring the average latency of nonauthoritative replies to queries in the last 60 minutes.
ibNetworkMonitorDNSNonAAT60AvgLatency (Integer)	Indicates the average latency in microseconds of nonauthoritative replies to queries in the last 60 minutes.
ibNetworkMonitorDNSNonAAT60Count (Integer)	Indicates the number of queries used to calculate the average latency of nonauthoritative replies in the last 60 minutes.
ibNetworkMonitorDNSNonAAT1440	File that contains the objects for monitoring the average latency of nonauthoritative replies to queries in the last 24 hours.
ibNetworkMonitorDNSNonAAT1440AvgLatency (Integer)	Indicates the average latency in microseconds of nonauthoritative replies to queries in the last 24 hours.
ibNetworkMonitorDNSNonAAT1440Count (Integer)	Indicates the number of queries used to calculate the average latency of nonauthoritative replies in the last 24 hours.

The table below describes the objects in `ibNetworkMonitorDNSAA`. You can send queries to retrieve values for these objects.

#### *ibNetworkMonitorDNSAA Objects*

<b>Object (Type)</b>	<b>Description</b>
ibNetworkMonitorDNSAAT1	File that contains the objects for monitoring the average latency of authoritative replies to queries in the last minute.
ibNetworkMonitorDNSAAT1AvgLatency (Integer)	Indicates the average latency in microseconds of authoritative replies to queries in the last minute.
ibNetworkMonitorDNSAAT1Count (Integer)	Indicates the number of queries used to calculate the average latency of authoritative replies in the last minute.
ibNetworkMonitorDNSAAT5	File that contains the objects for monitoring the average latency of authoritative replies to queries in the last five minutes.
ibNetworkMonitorDNSAAT5AvgLatency (Integer)	Indicates the average latency in microseconds of authoritative replies to queries in the last five minutes.
ibNetworkMonitorDNSAAT5Count (Integer)	Indicates the number of queries used to calculate the average latency of authoritative replies in the last five minutes.
ibNetworkMonitorDNSAAT15	File that contains the objects for monitoring the average latency of authoritative replies to queries in the last 15 minutes.
ibNetworkMonitorDNSAAT15AvgLatency (Integer)	Indicates the average latency in microseconds of authoritative replies to queries in the last 15 minutes.

ibNetworkMonitorDNSAAT15Count (Integer)	Indicates the number of queries used to calculate the average latency of authoritative replies in the last 15 minutes.
ibNetworkMonitorDNSAAT60	File that contains the objects for monitoring the average latency of authoritative replies to queries in the last 60 minutes.
ibNetworkMonitorDNSAAT60AvgLatency (Integer)	Indicates the average latency in microseconds of authoritative replies to queries in the last 60 minutes.
ibNetworkMonitorDNSAAT60Count (Integer)	Indicates the number of queries used to calculate the average latency of authoritative replies in the last 60 minutes.
ibNetworkMonitorDNSAAT1440	File that contains the objects for monitoring the average latency of authoritative replies to queries in the last 24 hours.
ibNetworkMonitorDNSAAT1440AvgLatency (Integer)	Indicates the average latency in microseconds of authoritative replies to queries in the last 24 hours.
ibNetworkMonitorDNSAAT1440Count (Integer)	Indicates the number of queries used to calculate the average latency of authoritative replies in the last 24 hours.

The below table describes the objects in `ibNetworkMonitorDNSSecurity`. When you enable the following, the SNMP traps with these objects are received:

- SNMP traps
- DNS network monitoring
- DNS alert monitoring

#### *ibNetworkMonitorDNSSecurityObjects*

Object (Type)	Description
ibNetworkMonitorDNSSecurityInvalidPort	Tracks the number of invalid DNS responses that arrive on invalid ports. For information about invalid ports, monitoring DNS transactions, see <a href="#">Monitoring Tools</a> . This object contains a subtree with six objects that track invalid ports within a certain time interval.
ibNetworkMonitorDNSSecurityInvalidTxid	Tracks the number of invalid TXIDs (DNS transaction IDs). For information about invalid TXIDs, monitoring DNS transactions, see <a href="#">Monitoring Tools</a> . This object contains a subtree with six objects that track invalid TXIDs within a certain time interval.
ibNetworkMonitorDNSSecurityInvalidPortOnly (Counter)	Tracks the number of DNS responses with both of the following conditions: <ul style="list-style-type: none"> <li>• Arrive on invalid ports</li> <li>• Have valid TXIDs</li> </ul>
ibNetworkMonitorDNSSecurityInvalidTxidOnly (Counter)	Tracks the number of DNS responses with both of the following conditions: <ul style="list-style-type: none"> <li>• Arrive on valid ports</li> <li>• Have Invalid TXIDs</li> </ul>
ibNetworkMonitorDNSSecurityInvalidPortCount (Counter)	Tracks the total number of invalid DNS responses that arrive on invalid ports.



ibNetworkMonitorDNSSecurityInvalidTxidCount (Counter)	Tracks the total number of DNS responses that have invalid DNS transaction IDs.
ibNetworkMonitorDNSSecurityInvalidTxidAndPort (Counter)	Tracks the number of DNS responses with both of the following conditions: <ul style="list-style-type: none"> <li>• Arrive on invalid ports</li> <li>• Have invalid TXIDs</li> </ul>

The following table describes the objects in `ibNetworkMonitorDNSSecurityInvalidPort`:

*ibNetworkMonitorDNSSecurityInvalidPort Objects*

Object (Type)	Description
ibNetworkMonitorDNSSecurityInvalidPort1 (Integer)	Tracks the number of invalid DNS responses that arrive on invalid ports in the last one minute.
ibNetworkMonitorDNSSecurityInvalidPort5 (Integer)	Tracks the number of invalid DNS responses that arrive on invalid ports in the last five minutes.
ibNetworkMonitorDNSSecurityInvalidPort15 (Integer)	Tracks the number of invalid DNS responses that arrive on invalid ports in the last 15 minutes.
ibNetworkMonitorDNSSecurityInvalidPort60 (Integer)	Tracks the number of invalid DNS responses that arrive on invalid ports in the last 60 minutes.
ibNetworkMonitorDNSSecurityInvalidPort1440 (Integer)	Tracks the number of invalid DNS responses that arrive on invalid ports in the last 24 hours.
ibNetworkMonitorDNSSecurityInvalidPortCount (Counter)	Tracks the total number of invalid DNS responses that arrive on invalid ports.

The following table below describes the objects in `ibNetworkMonitorDNSSecurityInvalidTxid`:

*ibNetworkMonitorDNSSecurityInvalidTxid Objects*

Object (Type)	Description
ibNetworkMonitorDNSSecurityInvalidTxid1 (Integer)	Tracks the number of DNS responses that have invalid DNS transaction IDs in the last one minute.
ibNetworkMonitorDNSSecurityInvalidTxid5 (Integer)	Tracks the number of DNS responses that have invalid DNS transaction IDs in the last five minutes.
ibNetworkMonitorDNSSecurityInvalidTxid15 (Integer)	Tracks the number of DNS responses that have invalid DNS transaction IDs in the last 15 minutes.

ibNetworkMonitorDNSSecurityInvalidTxid60 (Integer)	Tracks the number of DNS responses that have invalid DNS transaction IDs in the last 60 minutes.
ibNetworkMonitorDNSSecurityInvalidTxid1440 (Integer)	Tracks the number of DNS responses that have invalid DNS transaction IDs in the last 24 hours.
ibNetworkMonitorDNSSecurityInvalidTxidCount (Counter)	Tracks the total number of DNS responses that have invalid DNS transaction IDs.

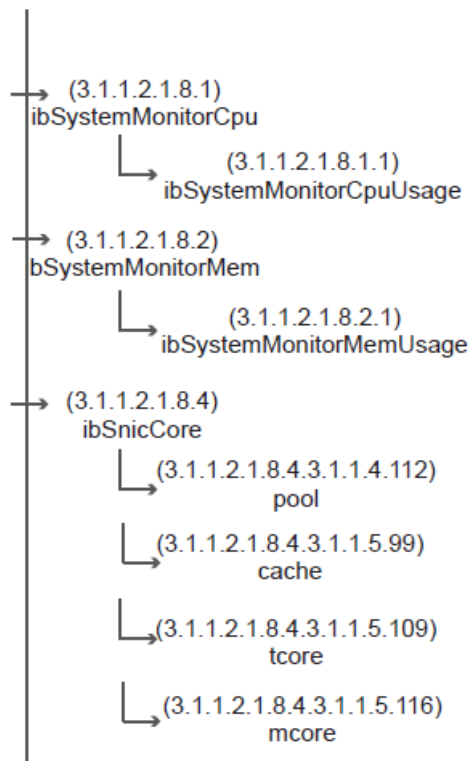
## ibSystemMonitor

As shown in the figure below, ibSystemMonitor (object ID 3.1.1.2.1.2.8) has the following subtrees:

- **ibSystemMonitorCpu:** Contains **ibSystemMonitorCpuUsage (Integer)** that reports the CPU usage of the appliance.
- **ibSystemMonitorMem:** Contains **ibSystemMonitorMemUsage (Integer)** that reports the memory usage of the appliance.
- **ibSystemMonitorSwap:** Contains **ibSystemMonitorSwapUsage (integer)** that reports the swap usage of the appliance.
- **ibSystemMonitorSnic:** Contains **ibSystemMonitorSnicStatsTable1 (integer)** and **ibSystemMonitorSnicStatsTable5 (integer)** that contains the smart NIC (SNIC) details. The OID 3.1.1.2.1.8.4.3.1.1.4.112.111.111.108 indicates that the member is a pool, 3.1.1.2.1.8.4.3.1.1.5.99.97.99.104.101 indicates that the member is a cache, 3.1.1.2.1.8.4.3.1.1.5.116.99.111.114.101 indicates that the cache is tcore, and 3.1.1.2.1.8.4.3.1.1.5.109.99.111.114.101 indicates that the cache is mcore.

### *ibSystemMonitor Objects3*

**(3.1.1.2.1.8)**  
**ibSystemMonitor**




**ibMemberServiceStatusTable**

As shown in the figure below, `ibMemberServiceStatusTable` (object ID 3.1.1.2.1.2.9) has one subtree, `ibMemberServiceStatusEntry`, which contains the following objects:

- `ibServiceName` (String) reports the names of the Infoblox services.
- `ibServiceStatus` (Integer) reports the status of the Infoblox services.
- `ibServiceDesc` (String) describes the details of the status.

`ibMemberServiceStatusTable` displays the current status of the Infoblox services on the appliance that you query. For an HA pair, this table displays the service status of the active node. If the appliance you query is the passive node of an HA pair, this table reflects the service status of the passive node, which can be "inactive" or "unknown." You can also query `ibMemberNodeServiceStatusTable` and `ibMemberPassiveNodeServiceStatusTable` that display system and hardware status on the queried appliance.

*ibMemberServiceStatusTable Objects*



### Infoblox Services for ibMemberServiceStatusTable

The following table lists the values and descriptions of the Infoblox services that appear in ibMemberServiceStatusTable:

#### *ibServiceName Values for ibMemberServiceStatus Table*

Value	Description	Definition
1	dhcp	DHCP service
2	dns	DNS service
3	ntp	NTP service
4	tftp	File distribution using the TFTP service
5	http-file-dist	File distribution using the HTTP service
6	ftp	File distribution using the FTP service
7	bloxtools-move	Moving the bloxTools service
8	bloxtools	The bloxTools environment

### Service Status

When you query the service status on an appliance, the response includes the status of the services. The below table shows the values and descriptions of the status. Note that for internal Grid operations, the NTP service is always in the "working" state even if it has been disabled through the Infoblox GUI.

### *ibServiceStates Values*

Value	Description	Definition
1	working	The service is functioning properly.
2	warning	The service is having some issues. Check the service or hardware function and the syslog to identify the problem.
3	failed	The service failed. Review the syslog to identify the problem.
4	inactive	The service is disabled or out of service.
5	unknown	The appliance cannot detect the current status of the service.

### *ibMemberNodeServiceStatusTable*

As shown in the below figure, *ibMemberNodeServiceStatusTable* (object ID 3.1.1.2.1.10) has one subtree, *ibMemberNodeServiceStatusEntry*, which contains the following objects:

- *ibMemberNodeServiceName* (String) reports the names of the system and hardware services.
- *ibMemberNodeServiceStatus* (Integer) reports the status of the services.
- *ibMemberNodeServiceDesc* (String) describes the details of the status.

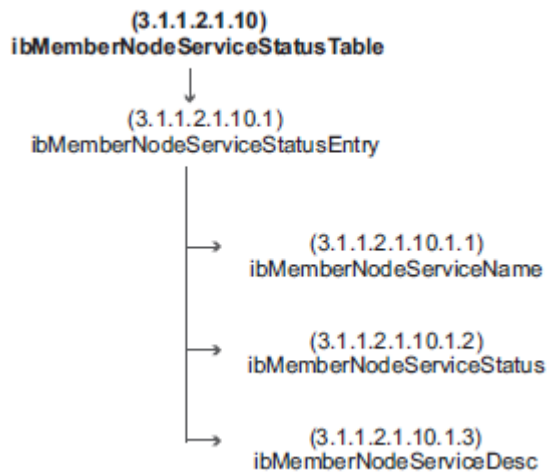
*ibMemberNodeServiceStatusTable* displays the current status of the system and hardware services on the appliance that you query. For example, when you query an independent appliance, this table shows the information about the independent appliance. When you query the VIP of an HA pair, this table shows the information about the active node. For the active node of the HA pair, you can also query *ibMemberPassiveNodeServiceStatusTable* to get the status of the passive node.



#### Note

For an independent appliance and the passive node of an HA pair, no information is returned when you query *ibMemberPassiveNodeServiceStatusTable*.

### *ibMemberNodeServiceStatusTable Objects*



### Infoblox Services for ibMemberNodeServiceStatusTable

The following table lists the values and descriptions of the Infoblox services that appear in ibMemberNodeServiceStatusTable:

#### *ibServiceName Values for ibMemberNodeServiceStatusTable*

Value	Description	Definition
9	node-status	Member status
10	disk-usage	Disk usage
11	enet-lan	LAN port
12	enet-lan2	LAN2 port
13	enet-ha	HA port
14	enet-mgmt	MGMT port
15	lcd	LCD
16	memory	Memory
17	replication	Replication service
18	db-object	Database capacity in %

19	raid-summary	RAID array
20	raid-disk1	RAID Disk 1 (For appliances with RAID arrays)
21	raid-disk2	RAID Disk 2 (For appliances with RAID arrays)
22	raid-disk3	RAID Disk 3 (For appliances with RAID arrays)
23	raid-disk4	RAID Disk 4 (For appliances with RAID arrays)
24	raid-disk5	RAID Disk 5 (For appliances with RAID arrays)
25	raid-disk6	RAID Disk 6 (For appliances with RAID arrays)
26	raid-disk7	RAID Disk 7 (For appliances with RAID arrays)
27	raid-disk8	RAID Disk 8 (For appliances with RAID arrays)
28	fan1	Fan 1
29	fan2	Fan 2
30	fan3	Fan 3
31	fan4	Fan 4
32	fan5	Fan 5
33	fan6	Fan 6
34	fan7	Fan 7
35	fan8	Fan 8
36	power-supply1	Power supply 1
37	power-supply2	Power supply 2
38	ntp-sync	NTP synchronization
39	cpu1-temp	CPU temperature
40	cpu2-temp	CPU temperature
41	sys-temp	System temperature

42	raid-battery	RAID battery
43	cpu-usage	CPU usage
44	ospf	OSPF
45	bgp	BGP
46	mgm-service	Multi-Grid management
47	subGrid-conn	Grid in Master Grid
48	network-capacity	Network capacity
49	reporting	Reporting service
50	dns-cache-acceleration	DNS Cache Acceleration services
51	ospf6	OSPF6
57	cloud-api	Cloud API service

### ibMemberPassiveNodeServiceStatusTable

As shown in the below figure, `ibMemberPassiveNodeServiceStatusTable` (object ID 3.1.1.2.1.2.11) has one subtree, `ibMemberPassiveNodeServiceStatusEntry`, which contains the following objects:

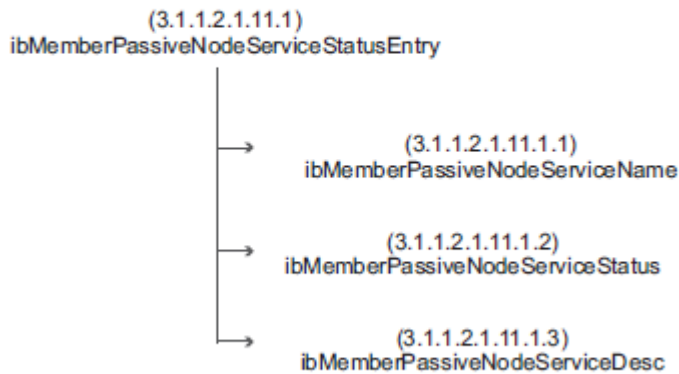
- `ibMemberPassiveNodeServiceName` (String) reports the names of the system and hardware services.
- `ibMemberPassiveNodeServiceStatus` (Integer) reports the status of the services.
- `ibMemberPassiveNodeServiceDesc` (String) describes details of the status.

`ibMemberPassiveNodeServiceStatusTable` displays the current status of the system and hardware services on the passive node of an HA pair when you query the VIP of the HA pair. For independent appliances and the passive nodes of HA pairs, this table does not display any status.

#### *ibMemberPassiveNodeServiceStatusTable Objects*



**(3.1.1.2.1.11)**  
**ibMemberPassiveNodeServiceStatusTable**



### ibDHCPOne MIB

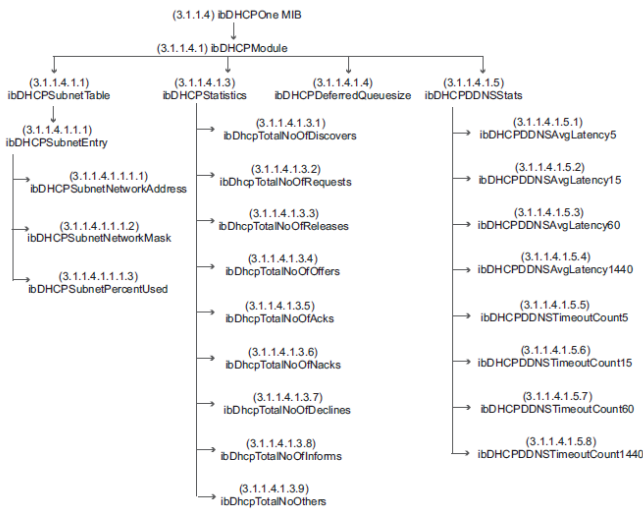
The ibDHCPOne MIB provides information about address usage within a subnet, DHCP lease statistics, and DHCP packet counts. It includes two modules, ibDHCPModule for IPv4 data and ibDHCPv6Module for IPv6 data.

### ibDHCPModule

The below figure illustrates the structure of the ibDHCPModule. (Note that the OIDs shown in the illustration do not include the prefix .1.3.6.1.4.1.7779.) ibDHCPModule contains the following objects:

- ibDHCPSubnetTable provides statistical data about the DHCP operations of the appliance.
- ibDHCPStatistics maintains counters for different types of packets.
- ibDHCPDeferredQueuesize tracks the total number of deferred DDNS updates that are currently in the queue to be retried. When DDNS updates are deferred due to timeout or server issues, the DHCP server puts these updates in this queue.
- ibDHCPDDNSStats monitors the average latency for the DDNS updates in microseconds and the number of timeouts during different time intervals.

### *ibDHCPModule*



ibDHCPSubnetTable provides statistical data about the DHCP operations of the appliance. It contains the following objects:

*ibDHCPSubnetTable*

Object (Type)	Description
ibDHCPSubnet Entry	File that contains the objects for monitoring DHCP operations on the appliance.
ibDHCPSubnetNetworkAddress (IbIpAddr)	The subnetworks, in IP address format, that have IP addresses for lease. A subnetwork may have many address ranges for lease.
ibDHCPSubnetNetworkMask (IbIpAddr)	The subnet mask in dotted decimal format.
ibDHCPSubnetPercentUsed (Integer)	<p>The percentage of dynamic DHCP addresses leased out at this time for each subnet. The percentage of dynamic DHCP addresses that are leased out are calculated as the [Total number of active leases in the range + Fixed addresses (leased or unused) + DHCP reservation addresses + Host addresses with DHCP enabled inside the specified subnet]/[Total number of IP address available in the range + Number of Fixed address inside the subnet + Number of IPv4 reservation addresses + Host addresses]</p> <p>For example:            Network: 10.192.16.0/24            Range: 10.192.16.10 – 10.192.16.20 (11 free IP addresses)            Fixed addresses: 10.192.16.3 (outside DHCP range) and 10.192.16.11 (within DHCP range)            There are 10 free IP addresses in the DHCP Range: 10.192.16.10 – 10.192.16.20 and one fixed address within the DHCP range</p> <p>IPv4 DHCP Utilization in Infoblox GUI: 16.6% (2/12) -&gt; For this case, (1 fixed address outside range + 1 fixed address within range) / (fixed address outside DHCP range + number of IP addresses in the DHCP range which includes the fixed address within range) * 100 = (2/12)*100 -&gt; 17%</p> <p>SNMP output from the CLI of the server:            Infoblox &gt; show snmp variable            .1.3.6.1.4.1.7779.3.1.1.4.1.1.3.11.49.48.46.49.57.50.46.49.54.46.48            IB-DHCPONE-MIB::ibDHCPSubnetPercentUsed."10.192.16.0" = INTEGER: 17</p>

Following is an example of the table as viewed through a MIB browser:

## MIBBrowserView1

ibDHCPSubnetTable		
ibDHCPSubnetNetworkAddress	ibDHCPSubnetNetworkMask	ibDHCPSubnetPercentUsed
1 10.0.0.0	255.0.0.0	0.0000
2 20.0.0.0	255.255.255.0	0.0000
3 20.0.1.0	255.255.255.0	0.0000
4 20.0.2.0	255.255.255.0	0.0000
5 20.0.3.0	255.255.255.0	0.0000
6 20.0.4.0	255.255.255.0	0.0000
7 20.0.5.0	255.255.255.0	0.0000
8 20.0.6.0	255.255.255.0	0.0000
9 20.0.7.0	255.255.255.0	0.0000
10 20.0.8.0	255.255.255.0	0.0000

## ibDHCPStatistics

ibDHCPStatistics maintains counters for different types of packets. The counters always start with zero when the DHCP service is restarted. Therefore, the numbers reflect the total number of packets received since the DHCP service was last restarted on the appliance. The ibDHCPStatistics module contains the following objects:

### ibDHCPStatistics

Object (Type)	Description
ibDhcpTotalNoOfDiscovers (Counter)	The number of DHCPDISCOVER messages that the appliance received. Clients broadcast DHCPDISCOVER messages when they need an IP address and network configuration information.
ibDhcpTotalNoOfRequests (Counter)	The number of DHCPREQUEST messages that the appliance received. A client sends a DHCPREQUEST message requesting configuration information, after it receives the DHCP OFFER message.
ibDhcpTotalNoOfReleases (Counter)	The number of DHCPRELEASE messages that the appliance received from its clients. A client sends a DHCP release when it terminates its lease on an IP address.
ibDhcpTotalNoOfOffers (Counter)	The number of DHCP OFFER messages that the appliance has sent to clients. The appliance sends a DHCP OFFER message to a client. It contains an IP address and configuration information.
ibDhcpTotalNoOfAcks (Counter)	The number of DHCPACK messages that the appliance sent to clients. It sends a DHCPACK message to a client to confirm that the IP address offered is still available.
ibDhcpTotalNoOfNacks (Counter)	The number of DHCPNACK messages that the appliance sent to clients. It sends a DHCPNACK message to withdraw its offer of an IP address.
ibDhcpTotalNoOfDeclines (Counter)	The number of DHCPDECLINE messages that the appliance received. A client sends a DHCPDECLINE message if it determines that an offered IP address is already in use.
ibDhcpTotalNoOfInforms (Counter)	The number of DHCPINFORM messages that the appliance received. A client sends a DHCPINFORM message when it has an IP address but needs information about the network.
ibDhcpTotalNoOfOthers (Counter)	The total number of DHCP messages other than those used in negotiation, such as DHCPFORCERENEW, DHCPKNOWN, and DHCPLEASEQUERY.

## ibDHCPDDNSStats

ibDHCPDDNSStats monitors the average latency for the DHCP DDNS updates in microseconds and the number of timeouts during different time intervals. The ibDHCPDDNSStats module contains the following objects:

### *ibDHCPStatistics*

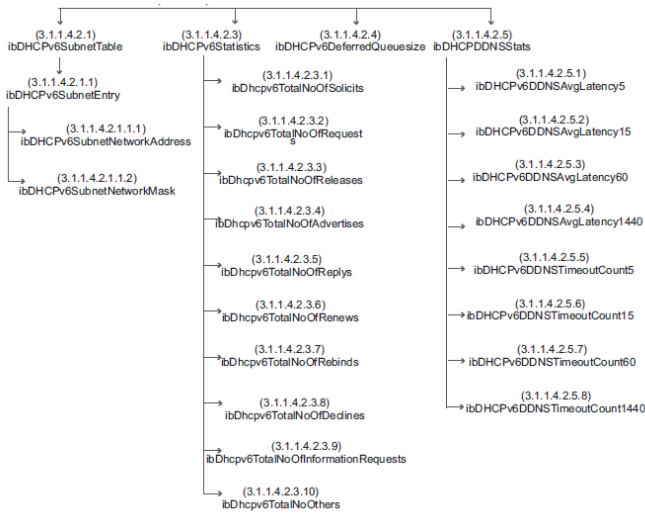
Object (Type)	Description
ibDHCPDDNSAvgLatency5 (Integer)	Indicates the average latency in microseconds of the DHCP DDNS updates in the last five minutes.
ibDHCPDDNSAvgLatency15 (Integer)	Indicates the average latency in microseconds of the DHCP DDNS updates in the last 15 minutes.
ibDHCPDDNSAvgLatency60 (Integer)	Indicates the average latency in microseconds of the DHCP DDNS updates in the last 60 minutes.
ibDHCPDDNSAvgLatency1440 (Integer)	Indicates the average latency in microseconds of the DHCP DDNS updates in the last 24 hours.
ibDHCPDDNSTimeoutCount5 (Integer)	The number of timeouts for the DHCP DDNS updates in the last five minutes.
ibDHCPDDNSTimeoutCount15 (Integer)	The number of timeouts for the DHCP DDNS updates in the last 15 minutes.
ibDHCPDDNSTimeoutCount60 (Integer)	The number of timeouts for the DHCP DDNS updates in the last 60 minutes.
ibDHCPDDNSTimeoutCount1440 (Integer)	The number of timeouts for the DHCP DDNS updates in the last 24 hours.

## ibDHCPv6Module

The figure below illustrates the structure of the ibDHCPv6Module, which contains the following objects:

- ibDHCPv6SubnetTable provides statistical data about the DHCPv6 operations of the appliance.
- ibDHCPv6Statistics maintains counters for different types of packets.
- ibDHCPv6DeferredQueuesize tracks the total number of deferred DDNS updates that are currently in the queue to be retried. When DDNS updates are deferred due to timeout or server issues, the DHCP server puts these updates in this queue.
- ibDHCPv6DDNSStats monitors the average latency for the DDNS updates in microseconds and the number of timeouts during different time intervals.

### *ibDHCPv6Module*



## ibDHCPv6SubnetTable

ibDHCPSubnetTable provides statistical data about the DHCPv6 operations of the appliance. It contains the following objects:

### *ibDHCPSubnetTable*

Object (Type)	Description
ibDHCPv6Subnet Entry	File that contains the objects for monitoring DHCPv6 operations on the appliance.
ibDHCPv6SubnetNetworkAddress (IbIpAddr)	The subnetworks, in IPv6 address format, that have IPv6 addresses for lease. A subnetwork may have many address ranges for lease.
ibDHCPv6SubnetNetworkMask (IbIpAddr)	The subnet mask in CIDR notation format.

## ibDHCPv6Statistics

ibDHCPv6Statistics maintains counters for different types of packets. The counters always start with zero when the DHCP service is restarted. Therefore, the numbers reflect the total number of packets received since the DHCP service was last restarted on the appliance. The ibDHCPv6Statistics module contains the following objects:

### *ibDHCPv6Statistics*

Object (Type)	Description
ibDhcpv6TotalNoOfSolicits (Counter)	The number of Solicit messages that the Grid member received, including Solicit messages embedded in Relay-Forward messages. A DHCP client sends a Solicit message to locate DHCP servers.

ibDhcpv6TotalNoOfRequests (Counter)	The number of Request messages that the Grid member received. A DHCP client sends a Request message to request one or more IP addresses and configuration parameters from a DHCP server.
ibDhcpv6TotalNoOfReleases (Counter)	The number of Release messages that the Grid member received. A DHCP client sends a Release message when it terminates its lease and releases its IP address.
ibDhcpv6TotalNoOfAdvertises (Counter)	The number of Advertise messages that the Grid member sent. When a DHCP server receives a Solicit message, it can respond with an Advertise message to indicate that the server is available for DHCP service.
ibDhcpv6TotalNoOfReplies (Counter)	The number of Reply messages that the Grid member sent. A DHCP server sends a Reply message that includes IP addresses and configuration parameters when it responds to Solicit, Request, Renew or Rebind message. It sends a Reply message with configuration parameters only when it responds to an Information-Request message.
ibDhcpv6TotalNoOfRenews (Counter)	The number of Renew messages that the Grid member received. A DHCP client sends a Renew message to a DHCP server to extend the lifetimes on the leases granted by the DHCP server and to update other properties.
ibDhcpv6TotalNoOfRebinds (Counter)	The number of Rebind messages that the Grid member received. A DHCP client sends a Rebind message to extend the lifetime of its lease and to update configuration parameters.
ibDhcpv6TotalNoOfDeclines (Counter)	The number of Decline messages that the Grid member received. A DHCP client sends a Decline message to a DHCP server when it discovers that the IP address offered by a DHCP server is already in use.
ibDhcpv6TotalNoOfInformationRequests (Counter)	The number of Information-Request messages that the Grid member received. A client sends an Information-Request message to retrieve configuration parameters, such as the IP addresses of DNS servers in the network.
ibDhcpv6TotalNoOfOthers (Counter)	The total number of DHCP messages other than those used in negotiation.

## ibDHCPv6DDNSStats

ibDHCPv6DDNSStats monitors the average latency for the DHCPv6 DDNS updates in microseconds and the number of timeouts during different time intervals. The ibDHCPv6DDNSStats module contains the following objects:  
*ibDHCPStatistics*

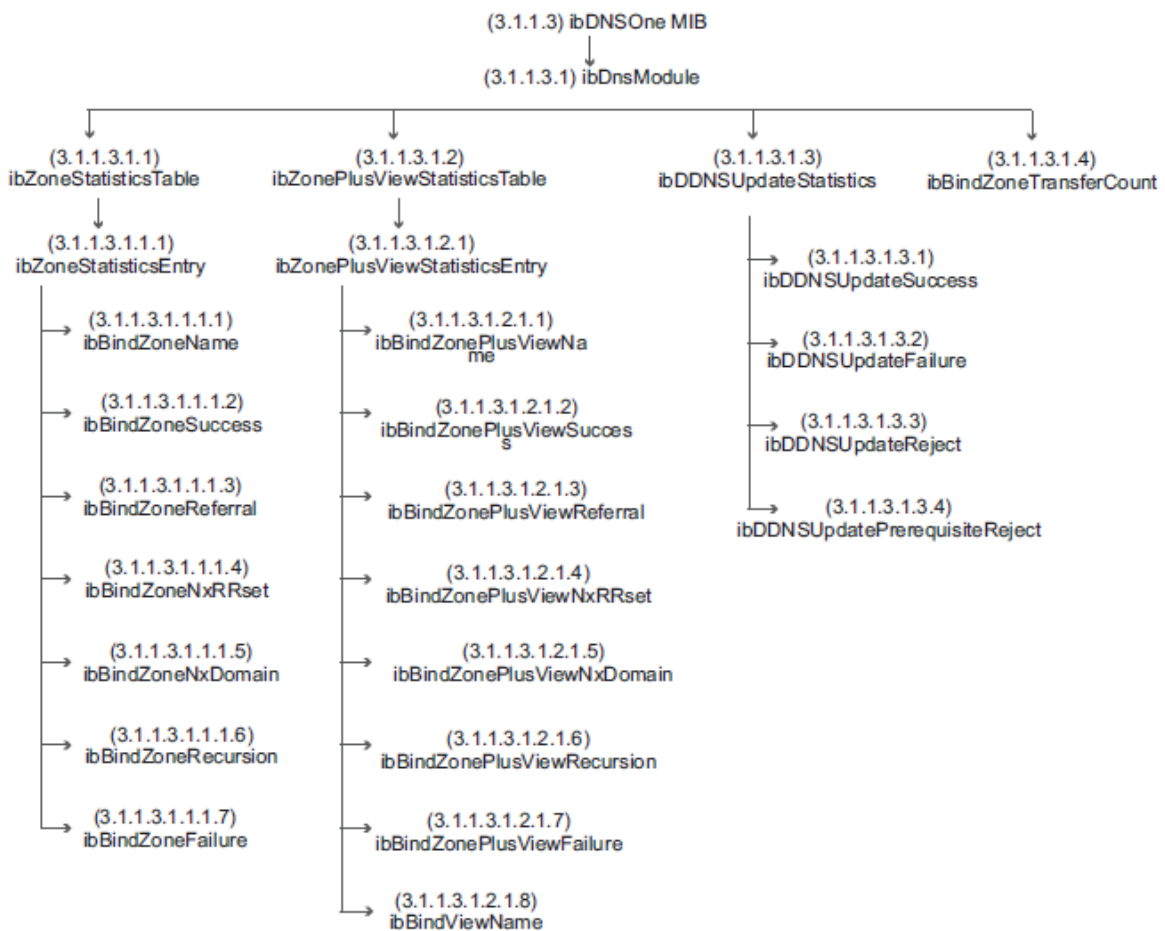
Object (Type)	Description
ibDHCPv6DDNSAvgLatency5 (Integer)	Indicates the average latency in microseconds of the DHCPv6 DDNS updates in the last five minutes.
ibDHCPv6DDNSAvgLatency15 (Integer)	Indicates the average latency in microseconds of the DHCPv6 DDNS updates in the last 15 minutes.
ibDHCPv6DDNSAvgLatency60 (Integer)	Indicates the average latency in microseconds of the DHCPv6 DDNS updates in the last 60 minutes.

ibDHCPv6DDNSAvgLatency1440 (Integer)	Indicates the average latency in microseconds of the DHCPv6 DDNS updates in the last 24 hours.
ibDHCPv6DDNSTimeoutCount5 (Integer)	The number of timeouts for the DHCPv6 DDNS updates in the last five minutes.
ibDHCPv6DDNSTimeoutCount15 (Integer)	The number of timeouts for the DHCPv6 DDNS updates in the last 15 minutes.
ibDHCPv6DDNSTimeoutCount60 (Integer)	The number of timeouts for the DHCPv6 DDNS updates in the last 60 minutes.
ibDHCPv6DDNSTimeoutCount1440 (Integer)	The number of timeouts for the DHCPv6 DDNS updates in the last 24 hours.

### ibDNSOne MIB

The ibDNSOne MIB provides DNS statistics about all zones in all views. The figure below illustrates the structure of the ibDNSOne MIB. (Note that the OIDs shown in the illustration do not include the prefix 1.3.6.1.4.1.7779.) The ibDNSOne MIB contains four subtrees: ibZoneStatisticsTable (Counter64), ibZonePlusViewStatisticsTable (Counter64), ibDDNSUpdateStatistics (Counter64), and ibBindZoneTransferCount (Counter64).

#### ibDNSOne MIB



## Using the DNS Zone Statistics Tables

`ibZoneStatisticsTable` and `ibZonePlusViewStatisticsTable` provide DNS statistics for all zones in all DNS views, including the default and all user-defined DNS views. You can use the information in these tables to calculate the total number of recursive queries on the DNS server. Depending on whether your DNS server is an authoritative or a caching-only server, you calculate the total number of recursive queries differently.

### `ibZoneStatisticsTable`

`ibZoneStatisticsTable` contains DNS statistics of all zones in the default DNS view. DNS statistics of user-defined DNS views are captured in `ibZonePlusViewStatisticsTable`.

`ibZoneStatisticsTable` includes a "summary" zone that provides global statistics for the DNS server, including statistics for all zones in the default and user-defined DNS views.

The syntax of the objects in `ibZoneStatisticsTable` uses a Counter64 format. In some cases, the counter format may not be compatible with SNMP toolkits that use a 32-bit counter. Ensure that you reconfigure or update these tools to use the Counter64 format. `ibZoneStatisticsTable` contains the following objects:

#### *ibZoneStatisticsTable*

Object (Type)	Description
<code>ibBindZoneName</code> (IbString)	DNS zone name. The index name for global statistics is "summary."
<code>ibBindZoneSuccess</code> (Counter64)	The number of successful responses since the DNS process started.
<code>ibBindZoneReferral</code> (Counter64)	The number of DNS referrals since the DNS process started.
<code>ibBindZoneNxRRset</code> (Counter64)	The number of DNS queries received for non-existent records.
<code>ibBindZoneNxDomain</code> (Counter64)	The number of DNS queries received for non-existent domains.
<code>ibBindZoneRecursion</code> (Counter64)	The number recursive queries received since the DNS process started.
<code>ibBindZoneFailure</code> (Counter64)	The number of failed queries since the DNS process started.

Following is an example of the table as viewed through a MIB browser:

#### *MIBBrowserView*



ibZoneStatistics Table						
Rotate            Refresh            Export <span style="float: right;">Poll</span>						
	ibBindZoneName	ibBindZoneSuccess	ibBindZoneReferral	ibBindZoneNxRRset	ibBindZoneNxDomain	ibBindZoneRecursion
1	parent	0	0	0	0	0
2	digzone	0	0	0	0	0
3	summary	1	3	0	0	0
4	ddnszone	0	0	0	0	0
5	reboottestzone	1	0	0	0	0
6	10.in-addr.arpa	0	0	0	0	0
7	20.in-addr.arpa	0	0	0	0	0
8	40.in-addr.arpa	0	0	0	0	0
9	zone-1-1-a.parent	0	0	0	0	0
10	zone-1-2-a.parent	0	0	0	0	0
11	zone-1-3-a.parent	0	0	0	0	0
12	zone-1-4-a.parent	0	0	0	0	0
13	zone-1-5-a.parent	0	0	0	0	0
14	zone-1-6-a.parent	0	0	0	0	0
15	zone-1-7-a.parent	0	0	0	0	0
16	zone-1-8-a.parent	0	0	0	0	0
17	zone-1-9-a.parent	0	0	0	0	0
18	zone-1-10-a.parent	0	0	0	0	0

### ibZonePlusViewStatisticsTable

ibZonePlusViewStatisticsTable provides DNS statistics about all zones in user-defined DNS views. DNS statistics about zones in the default view are captured in ibZoneStatisticsTable. Note that information in ibZonePlusViewStatisticsTable is rolled up to the "summary" zone in ibZoneStatisticsTable.

The syntax of the objects in ibZonePlusViewStatisticsTable uses a Counter64 format. In some cases, the counter format may not be compatible with SNMP toolkits that use a 32-bit counter. Ensure that you reconfigure or update these tools to use the Counter64 format. ibZonePlusViewStatisticsTable contains the following objects:

#### ibZonePlusViewStatistics Table

Object (Type)	Description
ibBindZonePlusViewName (IbString)	The zone name.
ibBindZonePlusViewSuccess (Counter64)	The number of successful responses since the DNS process started.
ibBindZonePlusViewReferral (Counter64)	The number of DNS referrals since the DNS process started.
ibBindZonePlusViewNxRRset (Counter64)	The number of DNS queries received for non-existent records.
ibBindZonePlusViewNxDomain (Counter64)	The number of DNS queries received for non-existent domains.
ibBindZonePlusViewRecursion (Counter64)	The number of queries that caused recursion since the DNS process started.
ibBindZonePlusViewFailure (Counter64)	The number of failed queries since the DNS process started.
ibBindViewName (IbString)	The DNS view name.

## Calculating Recursive DNS Queries

You can use the information in `ibZoneStatisticsTable` and `ibZonePlusViewStatisticsTable` to calculate the total number of recursive queries.

Following is an example of `ibZoneStatisticsTable` indexed by zone names in the default view:

```
index      ibBindZoneName ibBindZoneSuccess ibBindZoneReferral ibBindZoneNxRRset
            ibBindZoneNxDomain ibBindZoneRecursion ibBindZoneFailure
=====
"abc.com"  abc.com         0 0 0
           0 0 0
"summary"  summary        5 0 0
           0 0 0
"internal.com" internal.com    1 0 0
           0 0 0
```

Following is an example of `ibZonePlusViewStatisticsTable` indexed by zone names in all user-defined views:

```
index
ibBindZonePlusViewName ibBindZonePlusViewSuccess ibBindZonePlusViewReferral ibBindZonePlusViewNxRRset
            ibBindZonePlusViewNxDomain ibBindZonePlusViewRecursion ibBindZonePlusViewFailure
            ibBindViewName
=====
"ext1.com"  ext1.com       1 0 0 DNS1
           0 0 0
"ext2.com"  ext2.com       2 0 0 DNS1
           0 0 0
"ext3.com"  ext3.com       0 0 0 DNS2
           0 0 0
```

Use the `ibBindZoneSuccess` object in both tables to determine the total number of recursive queries. If your DNS server is a caching-only server, the total number of recursive queries is the number indicated in the `ibBindZoneSuccess` object of the "summary" zone. In this example, for a caching-only server, the total number of recursive queries is 5. If your DNS server is an authoritative server, add all the numbers in `ibBindZoneSuccess` for all zones in both tables, excluding the "summary" zone. In this example, the total is 4. You then subtract this number from the number in `ibBindZoneSuccess` of the "summary" zone. In this case, the total number of recursive queries is 1 for an authoritative DNS server.

## ibDDNSUpdateStatistics

ibDDNSUpdateStatistics provides statistical data about DDNS updates. The counters always start with zero when the DNS service is restarted. They report the total numbers since the DNS service was last restarted. ibDDNSUpdateStatistics contains the following objects:

### *ibDDNSUpdateStatistics*

Object (Type)	Description
ibDDNSUpdateSuccess (Counter64)	The number of successful dynamic DNS updates.
ibDDNSUpdateFailure (Counter64)	The number of all failed dynamic DNS updates, excluding those reported by the ibDDNSUpdateReject object.
ibDDNSUpdateReject (Counter64)	The number of dynamic DNS updates that failed because they were denied by the DNS server.
ibDDNSUpdatePrerequisiteReject (Counter64)	The number of dynamic DNS updates that failed because the prerequisites were not satisfied. This is also included in the total number of failures reported by the ibDDNSUpdateFailure object.

## ibBindZoneTransferCount

ibBindZoneTransferCount (Counter64) provides the total number of successful zone transfers from an Infoblox primary or secondary DNS server to a DNS client, since the DNS service was last restarted. Note that this counter tracks the number of successful full zone transfers (AXFRs) and incremental zone transfers (IXFRs).

## IB-DNSSERV-MIB

The IB-DNSSERV-MIB contains one object, ibDnsServConfig, which reports the DNS BIND version implemented by the NIOS software.

## IB-DNSHITRATIO-MIB

The IB-DNSHITRATIO-MIB contains one object, ibDnsHitRatio, which provides information about the DNS cache hit ratio. Note that the MIB variable (ibDNSOne) for cache hit rate has an OID of 1.3.6.1.4.1.7779.3.1.1.3.1.5.0, where 1.3.6.1.4.1.7779 is the prefix.

## IB-DNSQUERYRATE-MIB

The IB-DNSQUERYRATE-MIB contains one object, ibDnsQueryRate, which provides information about the DNS queries per second. Note that the MIB variable (ibDNSOne) for DNS query rate has an OID of 1.3.6.1.4.1.7779.3.1.1.3.1.6.0, where 1.3.6.1.4.1.7779 is the prefix.

## IB-DHCPSENV-MIB

The IB-DHCPSENV-MIB contains one object, ibDhcpv4ServerSystemDescr, which provides the DHCP server name and its DHCP version.

# Infoblox Reporting and Analytics

This section describes the Infoblox Reporting and Analytics solution and its features. It explains how to navigate through the user interface, view predefined dashboards, create personal reports and searches. It also provides best practices for customizing searches and setting permissions, and describes the reporting clustering feature and how to configure a reporting cluster.

It contains the following sections:

- [Setting Up Reporting and Analytics](#)
- [Supported Reporting Appliances and Storage Space](#)
- [Licensing Requirements](#)
- [Administrative Permissions for Reporting and Analytics](#)
- [Configuring an External Server for Search Result Exports](#)
- [Configuring Reporting Clustering](#)
- [Guidelines for Deploying Reporting Clusters](#)
- [Grid Reporting Properties](#)
- [Configuring IP Blocks and IP Block Groups](#)
- [Reporting User Interface Overview](#)
- [Predefined Dashboards](#)
- [Home Dashboards](#)
- [About Alerts](#)
- [About Reports](#)
- [About Dashboards](#)
- [Reporting Data Model](#)
- [Managing Reporting Data](#)

The Infoblox Reporting and Analytics solution automates the collection, analysis, and presentation of core network service data that assists you in planning and mitigating network outage risks so you can manage your networks more efficiently. It provides predefined dashboards and reports that capture useful information about the activities and performance of core network services. It also provides an enhanced reporting interface so you can create custom dashboards, reports, and alerts.



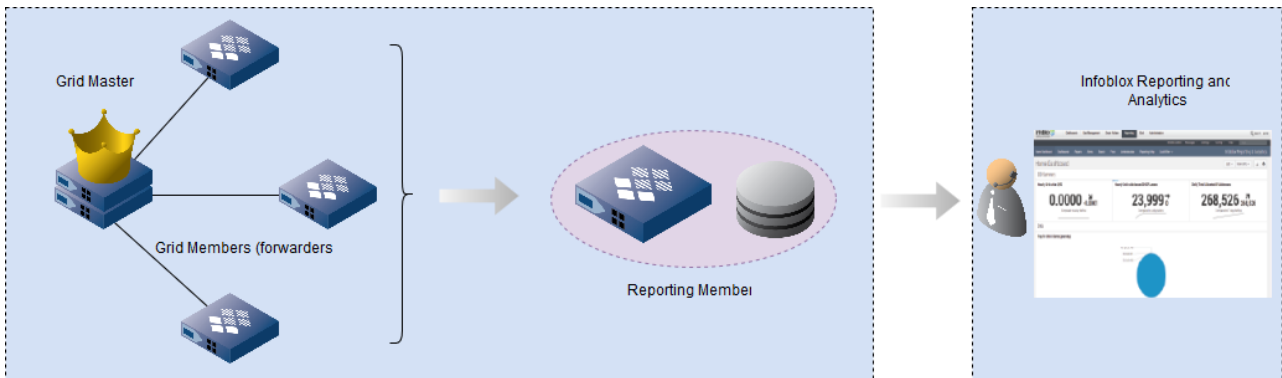
## Note

For Reporting and Analytics to function properly, ensure that you **DO NOT** create a SHA-256 4096 SSL key for the HTTPS certificate in your Grid because Java does not support SHA-256 with a 4096 key size.

Through reporting clustering, you can combine and configure multiple reporting members in a cluster. These reporting members work together to provide greater performance with higher data throughput and indexing capacity. The cluster also efficiently scales storage and indexing capacity. Reporting data is replicated among these reporting appliances to ensure continuous service even if one of the servers fails. You can configure more reporting appliances in multiple locations (sites) so that reporting data and service can be recovered from catastrophic disasters. Thus, the reporting clustering solution increases scale, offers higher reporting performance and greatly improves the reliability of the Reporting and Analytics solution. For information about this feature, see [About Reporting Clustering](#).

When you set up a reporting appliance with valid licenses in the Grid, the reporting server acts as an indexer that collects data from Grid members while the members are forwarders that transmit information to the reporting server. The reporting server indexes all raw data and transforms it into searchable events. Depending on your needs, you can enable certain Grid members as forwarders and disable others so the reporting server receives only the information you need from specific members. The following figure depicts the high-level configuration of the NIOS Reporting and Analytics solution:

*Infoblox Reporting and Analytics Overview*



The Infoblox reporting solution supports both IPv4 and IPv6 networks and you can configure a reporting member in IPv4, IPv6, or in dual mode (IPv4 and IPv6) network environment. An IPv4 reporting member uses IPv4 as the communication protocol, so you can add an IPv4 reporting member to an IPv4 or dual mode Grid. An IPv6 reporting member uses IPv6 as the Grid communication protocol, so you can add an IPv6 reporting member to an IPv6 or dual mode Grid. But a dual mode reporting member can use either IPv4 or IPv6 as the communication protocol, so you can add a dual mode reporting member to an IPv4, IPv6, or a dual mode Grid. For more information about how to set up the communication protocol, see [Changing the Communication Protocol for a Dual Mode Appliance](#).

## Related topic

[Guidelines for Upgrading the Reporting and Analytics Solution](#)

## Setting Up Reporting and Analytics

To enable and start using the Infoblox reporting solution on supported reporting appliances, complete the following:

1. Install a valid Reporting license on the reporting appliance, as described in [Licensing Requirements](#). Note that Infoblox now offer subscription licenses for the reporting solution.
2. Set administrative permissions, as described in [Administrative Permissions for Reporting and Analytics](#).
3. Configure reporting properties, as described in [Grid Reporting Properties](#).
4. Start the reporting service on the indexer and forwarders.

## Supported Reporting Appliances and Storage Space

Infoblox provides several reporting appliances to address different reporting requirements. The Report Categories, Related Data Sources, and Update Frequencies table lists the supported Trinzic Reporting platforms based on IP capacities and average DHCP leases and DNS queries per second, see [Configuring Reporting Clustering](#). There will be an impact on the performance of your reporting server when you perform many searches, download PDF reports, and send reports through emails and alerts. It is important to consider the reporting server configuration and make sure that it can handle the overall workload. A reporting appliance uses up to 95% of the total reporting disk space. The usable hard disk space on different appliance models is shown in the table below.

*Infoblox Reporting Appliances and their Usable Reporting Hard Disk Space*

Enterprise Model	Supported Infoblox Model	Reporting Hard Disk Space	Usable Reporting Hard Disk Space (GB)	Daily Maximum Data Consumption*
Trinzic Reporting 5005 Appliance (User-configurable)	IB-V5005	User-defined hard disk space	User-defined hard disk space	

Enterprise Model	Supported Infoblox Model	Reporting Hard Disk Space	Usable Reporting Hard Disk Space (GB)	Daily Maximum Data Consumption*
Very Large enterprises Service providers - Trinzic Reporting 4000 Appliance	IB-4000 (8x300GB HDD)	1555 GB	1477	20 GB
	IB-VM-4000	1500 GB	1477	20 GB
Large enterprises Service providers - Trinzic Reporting 2000 Appliance	IB-2205	2400 GB	1657	10 GB
	IB-2200-R2 (1000GB HDD)	1745 GB	1657	10 GB
	IB-2200-R1 (600GB HDD)	924 GB	877	10 GB
Mid-size Enterprises - Trinzic Reporting 1400 Appliance	IB-1405	1200 GB	1140	5 GB
	IB-1400-R2 (1000GB HDD)	833 GB	791	5 GB
	IB-1400-R1 (600GB HDD)	468 GB	444	5 GB
	IB-VM-1400	500 GB	475	5 GB
Mid-size Enterprises - Trinzic Reporting 800 Appliance	IB-805	1000 GB	950	5 GB
	IB-800	378 GB	359	2 GB
	IB-800	378 GB	359	1 GB
	IB-VM-800	250 GB	237	2 GB
	IB-VM-800	250 GB	237	1GB
	Mid-size Enterprises - Trinzic Reporting VM-800 (virtual appliance) Hyper-V**	IB-VM-800 Hyper-V**	250 GB	237
	IB-VM-800 Hyper-V**	250 GB	237	1 GB

 **Note**

The daily maximum data consumption includes all DNS, DDNS, IPAM, DHCP, Discovery, and system traffic or events from all members with data transmission enabled within the Grid. When data traffic exceeds the daily maximum, the reporting server sends an SNMP trap and email notification, if configured. After five (5) daily maximum warnings in a rolling period of 30 days, you cannot perform any reporting related functions. You must then contact Infoblox Technical Support to resolve the issue. Note that the reporting server continues to process incoming data during the violation state. However, you cannot view any reports or manage any reporting related functions until you fix the violation issue.

\*\*support for Reporting appliance models IB-VM-800, IB-VM-1400 starts from Microsoft Windows 2012 R2. Appliance model IB-VM-800 and IB-VM-1400 (Reporting) is supported for Microsoft Windows 2012 R2.

For information about the Trinzic Reporting platforms, their specifications, and how to install them as reporting appliances, refer to the respective installation guides, available on the Infoblox Support site.

## Licensing Requirements

You can install a valid Reporting license on a supported Trinzic Reporting platform and configure it as a reporting member solely for reporting purposes. You cannot add licenses to run other services, such as DNS and DHCP, on the reporting member. For information about Infoblox platforms that support reporting, see [Configuring Reporting Clustering](#).

Infoblox offers perpetual and subscription licensing programs to meet your business needs. You must obtain a new replacement license from Infoblox if you want to transfer an existing license to another member.

## Perpetual Licensing

A perpetual license is appliance-specific, meaning you can install only one perpetual license on each reporting appliance. There is no expiration date for this license type and you can use the Infoblox Reporting service for the lifetime of the appliance. However, you cannot install a perpetual license on an appliance that already has a valid DNS, DHCP, Microsoft Management, Query Redirection, or Multi-Grid Management license installed. You also cannot join a reporting member that has a perpetual license to a Grid that has a subscription license installed. A subscription license is a Grid-wide license and is not appliance-specific. Reporting members that have perpetual licenses can join a Grid that does not have a subscription license. For information about subscription licenses, see [Subscription Licensing](#) below.

### Note

All reporting appliances come with a 500 MB trial license. You can use the `set temp_license` CLI command to install the trial license. Available indexing capacity for a perpetual license: 1, 2, 5, 10 and 20 GB.

If you plan to configure reporting clustering, note that the overall indexing capacity for a reporting cluster increases based on the total licensing capacity for all reporting members. For example, if your Grid has three reporting members with 2 GB licensing capacity each, the overall cluster indexing capacity for the cluster becomes 6 GB. However, the stacking indexing capacity does not apply to trial licenses. In other words, if the three reporting members in your Grid all have trial licenses, the total cluster indexing capacity remains at 500 MB. For information about reporting clustering, see [About Reporting Clustering](#).

Note that you cannot install a perpetual reporting license on a Grid that has unrestricted reporting VMs as Grid members and vice versa. However, you can install a subscription license on a Grid that has unrestricted reporting VMs as Grid members. Unrestricted reporting VMs are reporting appliances for which you can arbitrarily configure CPU, memory, and disk size. You can arbitrarily configure CPU, memory, and disk size on IB-V5005 appliance.

## Subscription Licensing

A subscription license is a Grid-wide license that you install on the Grid Master, and it is not appliance-specific. A subscription license can either be permanent or have an expiration date and Grid Manager (the Infoblox GUI) displays a 90-day and 30-day warning prior to the actual expiration date. The effective start date for a subscription license starts on the day you generate the license. Note that a reporting member that has a perpetual license is not allowed to join the

Grid that has a subscription license.

In a dynamic license pool, expiration warnings are based upon the expiration dates of the sub pools. When any licenses in a sub pool expire, the total available licenses for a pool drops accordingly. On the other hand, a perpetual static license allocated to an appliance does not expire. The appliance continues to run the feature for the lifetime of that appliance. For more information about Managing Grid-wide Licenses, see [Managing Licenses](#).

 **Note**

Available indexing capacity for a subscription license: 1, 2, 5, 10, 20, 50, 100, 200, and 500 GB. Contact your Infoblox representative for pricing and availability information.

## About License Violations

License violations occur when the reporting appliance exceeds the maximum allowed daily volume or when the indexed data type is not allowed for the license type. There might be an impact on the reporting service performance if violations occur. License violation is reported in the **Reporting** tab of Grid Manager. During the license violation period, the reporting server continues to index data but the search function will cease to be operational. This means that the **Reporting** tab is available with the following warning message displayed: "Reporting Service is unavailable." You can apply a **Reset** license to remove the violations. Contact Infoblox Technical Support to obtain a **Reporting Reset** license.

Note the following behavior for such violations:

- NIOS continues to index data; however, you will not be able to use the search feature.
- You can use the reporting search when the number of violations in the previous 30 days is within the limit. To obtain valid licenses, contact your Infoblox representative or Infoblox Technical Support.
- If there are five consecutive violations in five consecutive days, then the search feature is disabled for the next 25 days until a **Reset** license is installed. If necessary, contact Infoblox Technical Support to obtain a **ReportingReset** license so you can reset the current license violations.
- During a license violation, summary indexing does not work and reports that use the summary index will not display output. For information about reports that use the summary index, see [Reporting Data Model](#).

## Administrative Permissions for Reporting and Analytics

You must have the appropriate permissions to access, view, edit, and clone searches, dashboards, reports, and alerts that are available in the **Reporting** tab. NIOS synchronizes all users and admin groups as users and roles on the reporting server. In NIOS, admin groups may have the setting "superusers" enabled or disabled. NIOS groups that have "superusers" enabled will correspond to roles with administrative permissions on the reporting server. This documentation will refer to this role as a superuser role and the associated users as superusers. NIOS groups that do not have "superusers" enabled will not have administrative permissions on the reporting server. This role will be referred to as a non-superuser role and the associated users are limited-access users. The default NIOS admin group "admin-group" will have a superuser role on the reporting server called "infoblox-admin-role." The default NIOS admin named "admin" will have a user name on the reporting server called "infoblox-admin."

 **Note**

When you enable the reporting service, the *splunk-reporting-group* is created automatically for authentication by the Reporting and Analytics App on the Grid. It is an internal admin group, so you must not add users to this group or rename it.

By default, superusers have full access to the Reporting and Analytics App and limited-access users do not. Superusers can grant permissions to limited-access users so they can access the Reporting and Analytics App. Limited-access users will not have access to the **Reporting** tab until superusers grant them this access. For information about Splunk roles, refer to the Splunk documentation. Note that the reporting service does not support non-ASCII characters in the names of admin groups and admin users. When you include ASCII characters and space in the admin user name, make sure that its length does not exceed 33 characters.



When configuring NIOS, only superuser and users with the Grid Reporting Properties permission can configure Grid Reporting Properties. Limited-access users can view the Grid Reporting Properties. When superusers create or delete any admin groups in NIOS, the corresponding roles are created or deleted on the reporting server. You cannot directly create or delete roles on the reporting server. If superuser permission is granted or removed from a NIOS group, this permission change will be reflected in the corresponding role on the reporting server. When you modify the name of an admin group on NIOS, make sure that you grant permissions again for all corresponding roles that are created on the Reporting and Analytics App for this renamed admin group. For information about granting permissions, see [Granting Permissions](#) below.



#### Note

Limited-access users cannot create, modify, and delete any user role for the Splunk App.

## Granting Permissions

Superusers can grant App permissions to limited-access users. Permissions for all reporting objects are migrated to the new Reporting solution and managed through the new user interface after an upgrade. You may see some new built-in roles, such as **Everyone**, when configuring Reporting permissions. For best practices, do not alter permissions for these new built-in roles. In addition, Reporting Dashboard and Reporting Search global permissions have been removed. If an admin group or admin role was granted these permissions before an upgrade, the permissions will still be displayed after an upgrade. However, they won't take any effect. The Grid Reporting Properties permission is retained.



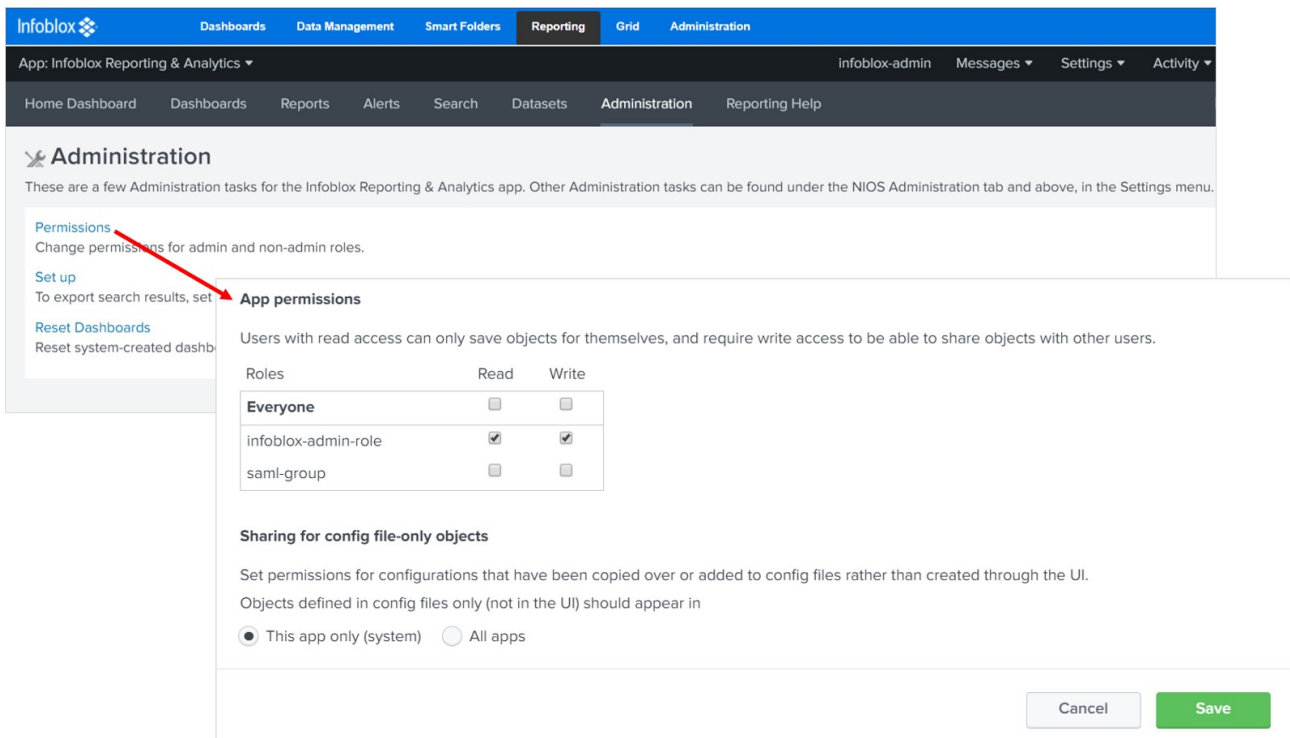
#### Note

The **App Permission** and **Edit Permissions** sections display only your assigned admin role on the reporting server. Therefore, you can only view the admin role to which your user name was assigned. Only the default system admin can view all admin roles.

To set permissions:

1. From the **Reporting** tab -> select the **Administration** tab.
2. Click **Permissions** to manage permissions.
3. In the **App permissions** dialog box, set permissions to **Read** and/or **Write** for the roles listed. You must have at least one permission (Read or Write) to execute a task on the **Reporting** tab.
4. Click **Save**.

*Setting Permissions for Reporting and Analytics App*



## Editing Permissions

Users can edit permissions for objects where they are the owner, such as dashboards, reports, and alerts. When a user creates a new report, dashboard or an alert, it is only available to that user. To make that object available to other users, you can do the following (only if your permissions allow you to do so):

- Make an object available to all users.
- Restrict or expand access to all or specific objects by roles.
- Set Read or Write permissions at the Reporting level for roles. For information about users and roles, refer to the Splunk documentation.

To modify permissions:

1. From the **Reporting** tab -> select **Dashboards** or **Reports** or **Alerts**.
2. From the **Edit** drop-down list, select **Edit Permissions**.
3. Specify the following:
  - Display for **Owner**, **App**, or **All Apps**. For more information about **All Apps** permissions, refer to the Splunk documentation.
  - Read and write privileges for users. By default, the permissions are set on the object. You cannot remove the **Read** and **Write** permissions completely. However, you can change these permissions from **Read** to **Write** or vice versa, based on your permissions. For more information about permissions, refer to the Splunk documentation.
4. Click **Save**.

## Configuring an External Server for Search Result Exports

You can configure an FTP, SCP, or TFTP server to which you plan to schedule the export of search results. Only superusers can configure the remote server and limited-access users cannot do so. When you upgrade to 7.3.x, make sure that you reconfigure an FTP, SCP, or TFTP server, even if you have already scheduled the export of search results. The **Set up** page to configure an FTP, SCP, or TFTP server is displayed for all new installations and upgrades if you have not previously scheduled the export of search results. If you have configured any scheduled export of search results in a

previous NIOS release, NIOS migrates one of the server settings to the Reporting and Analytics App. In this case, the **Set up** page is not displayed. However, you can still access the **Set up** page from the **Reporting** tab > **Administration** tab, as illustrated in Setup Page to specify Server for Exporting Search Results figure below.

To configure an FTP, SCP, or TFTP server:

1. From the **Reporting** tab -> select the **Administration** tab -> click **Set up**.
2. To configure a remote server, complete the following:
  - **Username:** Enter the username of your server account.
  - **Password:** Enter the password of your server account.
  - **Confirm Password:** Enter the same password.
  - **Protocol:** Select **FTP** or **SCP** or **TFTP** from the drop-down list.
  - **Host/IP Address:** Enter the host IP address.
  - **Host Port:** Enter the port number on the selected server.
  - **Destination Path:** Enter the path and the file name of the export file. For example, you can enter /export/Infoblox\_2009\_10\_20\_15\_30 on a Linux server, or c:\export\Infoblox\_2009\_10\_20\_15\_30 on a Microsoft Windows server.
3. Click **Save**.

### Setup Page to specify Server for Exporting Search Results

The screenshot shows the Infoblox Reporting & Analytics Administration interface. The top navigation bar includes 'Dashboards', 'Data Management', 'Smart Folders', 'Reporting', 'Grid', and 'Administration'. The 'Reporting' tab is active. Below the navigation bar, there are links for 'Home Dashboard', 'Dashboards', 'Reports', 'Alerts', 'Search', 'Datasets', 'Administration', and 'Reporting Help'. The main content area is titled 'Administration' and contains a list of tasks: 'Permissions', 'Set up', and 'Reset Dashboards'. A red arrow points from the 'Set up' link to the 'Set Up File Server for Exporting Search Results' form. The form is titled 'Enter File Server Settings' and contains the following fields: Username (admin), Password (empty), Confirm password (empty), Protocol (SCP), Host/IP Address (10.120.21.222), Host Port (22), and Destination Path (/import/qaddi/admin). At the bottom right of the form are 'Cancel' and 'Save' buttons.

## Configuring Reporting Clustering

You can add higher scale, performance, and reliability to the Reporting and Analytics solution by using the reporting clustering feature. Through reporting clustering, you can combine and configure multiple reporting members in a cluster.

The reporting clustering feature offers high availability and disaster recovery. You can configure one or more reporting appliances in multiple locations (sites). Reporting data is replicated among these reporting appliances to ensure redundancies and continuous service even if one of the servers fails. For example, if one of the reporting members has

operational issues, reports and dashboards will use backup copies of the data on other reporting members in the cluster to ensure continuous reporting service. When a new reporting member joins the cluster, you do not need to reconfigure and restart your forwarders to send data to the new reporting member as the Grid Master automatically notifies all forwarders about the new member. In addition, data indexed on the new reporting member participates in searches that support reports and dashboards. Thus, the new reporting clustering solution increases scale, offers higher reporting performance and greatly improves the reliability of the Reporting and Analytics solution.

The reporting clustering feature is supported on appliances running NIOS version 7.3.200 and later. For new installations and upgrades from a previous NIOS release to NIOS 7.3.200 and later, the appliance is set to the single indexer mode by default. For more information about how reporting clustering works and the types of clustering mode, see [Clustering Overview](#) and [Reporting Cluster Modes](#) below.

 **Note**

Reporting clusters are not supported in a Multi-Grid configuration.

## Clustering Overview

The concept of reporting clustering is to set up a group of reporting members within one site (location) or across multiple sites. When you configure multiple reporting members within one site, you are setting up a single-site cluster. Configuring multiple reporting members across different sites gives you a multi-site cluster, as illustrated in the figure [Sample Multi-Site Reporting Cluster](#). Single-site clusters and multi-site clusters below, offer the benefits of high availability and disaster recovery. Without reporting clustering, a reporting member is known as a single indexer.

A reporting cluster, either single-site or multi-site, consists of the following components that work together to perform reporting and clustering activities:

- **Cluster Master:** The cluster master coordinates all clustering activities and always runs on the Grid Master.
- **Indexer** (also known as cluster peer): An indexer that collects, processes, and indexes reporting data. It can also function as the originating indexer (source peer) or a replication target (target peer).
- **Forwarder:** A forwarder sends reporting data to the indexer for processing.
- **Search Head:** A search head handles search queries and distributes search requests to indexers in the cluster. One of the reporting members in the cluster will have double duties of being the indexer and search head.
- **Replication Factor:** This factor defines the number of copies of reporting data the cluster replicates and maintains. This is set to 2 by default for both single-site and multi-site clusters so the clusters can tolerate one reporting member failure without losing any data (since there will still be another copy of data available in the cluster).
- **Search Factor:** This factor defines the number of copies of searchable data. This is set to 2 for single-site clusters and set to 1 for multi-site clusters so the cluster can tolerate one reporting member failure without impacting search results (since there will still be a searchable copy of data available in the cluster).
- **ReportingSite Extensible Attribute:** This is an extensible attribute that you associate with reporting members in a multi-site cluster. For more information, see [ReportingSite Extensible Attribute](#) below.

In a Grid that includes a reporting cluster, the Grid Master coordinates various activities across reporting members. In the reporting cluster, a reporting member can act as an indexer and/or a search head. It also participates in peer-to-peer data replication.

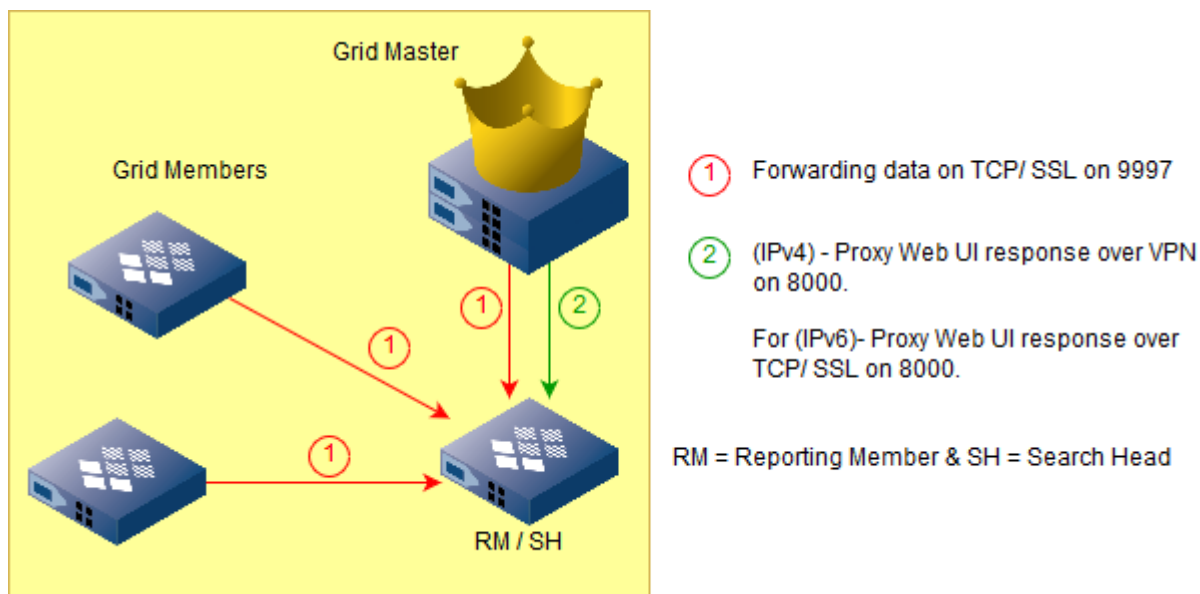
To configure a reporting cluster, you must first set up all the reporting members and enable the reporting service in the Grid. You can then select the reporting clustering type for your cluster. For more information about cluster types, see [Reporting Cluster Modes](#) below. When you configure the reporting cluster, you must use an NTP server to synchronize the time of the Grid Master, Grid members and reporting members.

## Reporting Cluster Modes

You must first enable the reporting service and configure one or more reporting members as needed before configuring the reporting cluster. When you enable reporting clustering, the Grid Master, forwarders, and reporting members use specific ports for network communication. [Ports Required for IPv4 and IPv6 Single Indexer](#), [Port Requirement for IPv4 and IPv6 Single-site Clustering](#), and [Port Requirement for IPv4 and IPv6 Multi-site Clustering](#) figures below illustrates whether the network communication is over TCP/SSL or VPN, and ports that you can use for the reporting service.

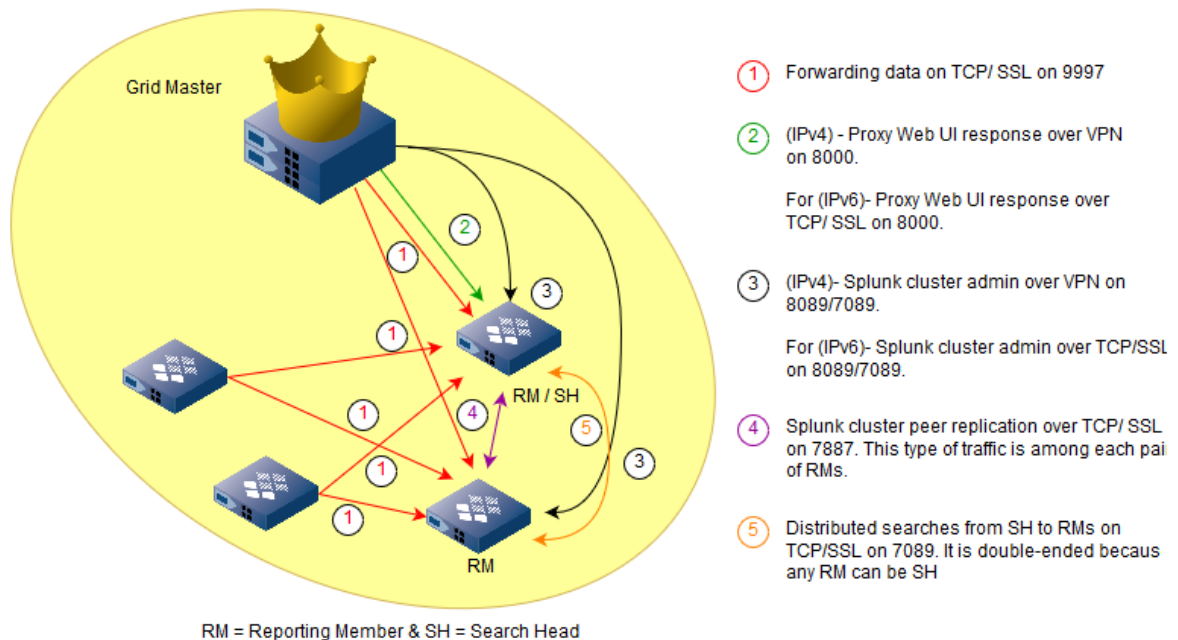
- **Single Indexer:** This is the traditional configuration that works on one reporting server (indexer). The forwarder sends reporting data to the indexer and the indexer indexes the data. This is the default configuration when you enable the reporting service for new installations and when you upgrade from a previous NIOS release to NIOS 7.3.200 and later.

*Ports Required for IPv4 and IPv6 Single Indexer*



- **Single-Site Cluster:** In a single-site cluster, the Grid Master is also the cluster master and all reporting members are cluster indexer peers. NIOS selects a peer and configures it as the search head to handle search queries. If the selected search head goes down, NIOS automatically selects another search head among the reporting members in the same site. All other Grid members (non-reporting members) are considered forwarders that send reporting data to the cluster peers for processing. You must configure at least two reporting members that are located in the same site (location). By default, the replication factor and search factor for a single-site cluster are set to 2. Note that you can upgrade your configuration from a single-site cluster to a multi-site cluster. However, once configured, you cannot change your configuration back to a single indexer. For information about how to configure a single-site cluster, see *Configuring Reporting Clusters* below.

*Port Requirement for IPv4 and IPv6 Single-site Clustering*



- **Multi-Site Cluster** - A multi-site clustering configuration is useful when you want to manage multiple reporting sites at different locations, with each site having its own set of indexers. The multi-site clustering configuration is valid only when you associate all the reporting members in the cluster with the predefined **ReportingSite** extensible attribute. For information about the ReportingSite extensible attribute, see ReportingSite Extensible Attribute below. In a multi-site cluster, you configure one of the sites as the primary site, and then plan other sites in a specific order. This order defines the next site of indexers to which the forwarders send data when the primary site is out of service. Note that all Grid members send data only to indexers in the primary site. You can designate a new primary site either by using the *Grid Reporting Properties* editor, or using the `set promote_master` CLI command. For more information about the CLI command, refer to the *Infoblox CLI Guide*. A multi-site cluster must have at least two sites with two reporting members in each site, as illustrated in the Sample Multi-Site Reporting Cluster figure. The first reporting site that you configure is the primary site, which also hosts the search head for the cluster. If the search head goes down, the Grid Master automatically chooses an available reporting member in the same site as the search head. If all the indexers in a site go down, or if you want to change the search head to another site, then you must manually redefine the primary site. Note that you must make one of the active sites as the primary site. In a multi-site cluster, the search factor (also known as the site search factor) determines both the number of searchable copies that the entire cluster maintains and the number of copies that each site maintains. By default, the search factor is set to 1 and the replication factor is 2 in a multi-site cluster.

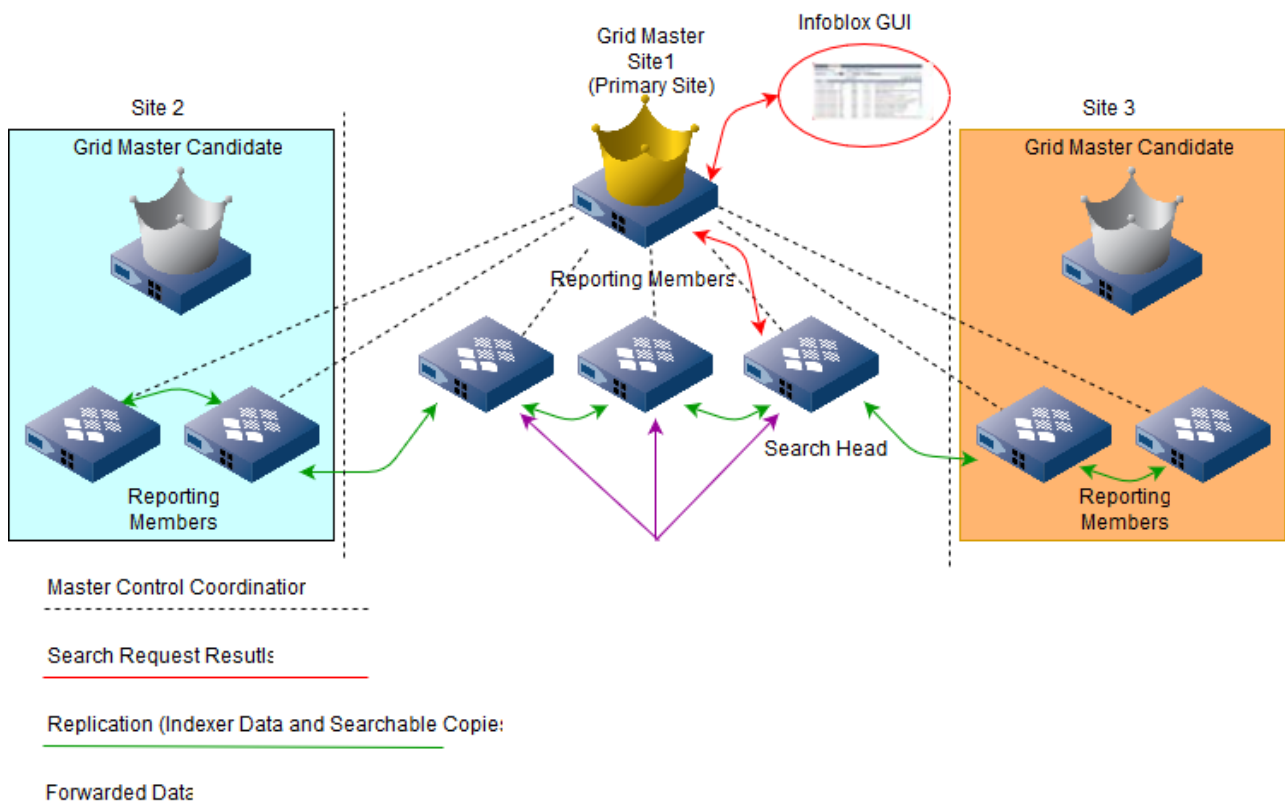
**Note**

You can change your configuration from a single indexer to a single-site cluster or multi-site cluster and from a single-site cluster to a multi-site cluster. However, you cannot revert your configuration from a multi-site cluster to a single-site cluster or to a single indexer.

## Clustering Data Replication

When you change the configuration from a single indexer to a single-site cluster or multi-site cluster and from a single-site cluster to a multi-site cluster, the replication of data will start only for the new data that are created after you have completed the cluster mode configuration. When you change the configuration, the replication of new data starts only after you have completed the clustering configuration. Any data created prior to switching are restored on the primary site and are not replicated on the secondary site. To manage your reporting clustering data efficiently, see [Guidelines for Deploying Reporting Clusters](#).

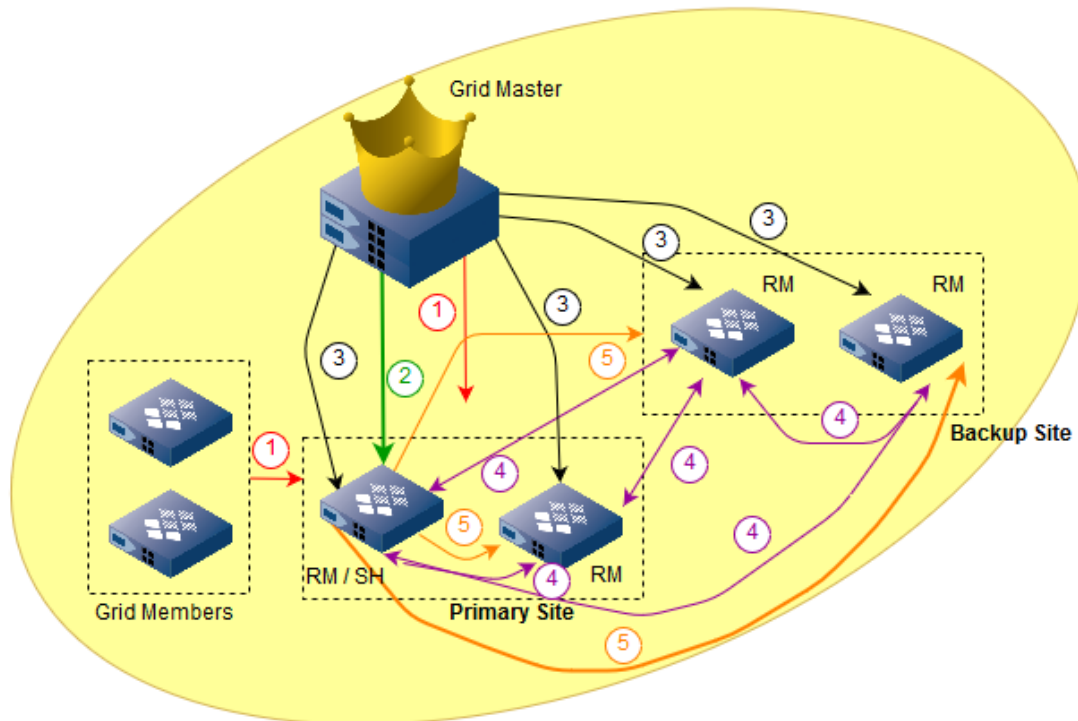
### Sample Multi-Site Reporting Cluster



For more information about how reporting cluster works, refer to the Splunk documentation at <https://docs.splunk.com/Documentation/Splunk/8.2.4/Indexer/Basicclusterarchitecture>.

### Port Requirement for IPv4 and IPv6 Multi-site Clustering





- ① Forwarding data on TCP/ SSL on 9997
  - ② (IPv4) - Proxy Web UI response over VPN on 8000.  
For (IPv6)- Proxy Web UI response over TCP/ SSL on 8000.
  - ③ (IPv4)- Splunk cluster admin over VPN on 8089/7089.  
For (IPv6)- Splunk cluster admin over TCP/SSL on 8089/7089.
  - ④ Splunk cluster peer replication over TCP/ SSL on 7887. This type of traffic is among each pair of RMs.
  - ⑤ Distributed searches from SH to RMs on TCP/SSL on 7089. It is double-ended because any RM can be SH
- RM = Reporting Member & SH = Search Head

## ReportingSite Extensible Attribute

NIOS defines the ReportingSite extensible attribute for use by the multi-site cluster reporting configuration. You must associate a ReportingSite extensible attribute value with the reporting members defined in the cluster. For more information, see [Assigning a Reporting Site EA Value to a Multi-Site Cluster](#) below. Note that your multi-site cluster configuration is invalid unless you assign ReportingSite values to all the reporting members that are part of the cluster. You can add up to five ReportingSite extensible attributes, and view and edit the ReportingSite extensible attributes in the **Administration** tab -> **Extensible Attributes** tab in Grid Manager. You can view the **ReportingSite** extensible attribute values in the **Grid** -> **Grid Manager** -> **Members** tab in Grid Manager. The **ReportingSite** column is not available if you customize the Results table. You can enable the **ReportingSite** column by selecting **Columns** -> **Edit Columns**. For information about customizing tables, see [About the Grid Manager Interface](#). You can also use the **Group Results** function to group reporting members that contain the same ReportingSite extensible attributes. For information about grouping members by extensible attributes, as described in [Grouping Members by Extensible Attribute](#), see [Adding Grid Members](#).



As illustrated in Sample Multi-Site Reporting Cluster above, the ReportingSite value "site1" is assigned to a site within a multi-site cluster and to the reporting members RM1, RM2 and RM3. The ReportingSite value "site2" is assigned to a different site in the cluster and to reporting members RM4, RM5 and RM6. If the search head goes down, the Grid Master automatically chooses an available reporting member in the same site to be the search head.

 **Note**

When you modify the ReportingSite extensible attribute value for any indexers in a multi-site cluster, ensure that you validate the configuration again, as described in [Validating Reporting Clustering Configuration](#) below.

## Monitoring Reporting Cluster Status

After you have set up reporting members and defined clustering type, you can monitor the cluster status through the following:

- View the reporting member service status, as described in Monitoring Grid Services, see [Monitoring Services](#).
- Check license usage by the reporting member, as described in [Home Dashboards](#).

## Promoting the Grid Master Candidate in Multi-Site Clustering

If the Grid Master fails and all other reporting members are up and running in a multi-site reporting cluster, you must promote the Grid Master Candidate to the Grid Master by using the CLI command `set promote_master`. For information about the CLI command, refer to the [Infoblox CLI Guide](#).

## Reporting Categories and Related Data Sources

The reporting member uses two types of data sources to generate reports: file-based data sources and script-based data sources. When the reporting member is down or unreachable, file-based data sources are queued until the reporting member is up and running. However, the script-based data sources are lost if the size of the queued data exceeds 500 KB.

The amount of data in the queue is managed as follows:

- Rotates the reporting syslog files (extracted from `/var/log/messages`) at 120 MB retaining one older file. The data in the queue depends on the file size when the reporting member becomes unreachable.
- The CSV files overwrite the oldest data with the new data at regular intervals. So, the CSV file contains only the latest events.

 **Note**

If you desire older information to be kept, use any of the export methods to daily export this data to a file.

The table below lists the reports provided by the reporting server, report categories, and the source type, data source type (file or script-based), and queue data update frequencies for each report:

### *Report Categories, Related Data Sources, and Update Frequencies*

Report Category	Reports	Source Type	Data Source (file-based or scriptbased)	Update Frequency
Device	Inactive IP Addresses	ib:reserved 2	file-based (syslog)	Rotates at 120 MB; retains one older copy; queued data is between 120 MB and 240 MB

Report Category	Reports	Source Type	Data Source (file-based or scriptbased)	Update Frequency
	Port Capacity Utilization by Device Port Capacity Trend Port Capacity Delta by Device	ib:reserved 2	file-based (csv)	Overwritten every 6 hours
	End Host History	ib:discovery:end_host_activity	file-based (csv)	Overwritten every 24 hours
	IP Address Inventory			
	Network Inventory			
	IPAMv4 Device Networks			
	Device Interface Inventory			
	Device Inventory	ib:reserved 2	file-based (csv)	Overwritten every 24 hours
	Device Components			
	Device Advisor	ib:reserved 2	file-based (csv)	Overwritten every 24 hours
DHCP Performance	DHCP Message Rate Trend	ib:dhcp:message	file-based (csv)	Overwritten every 1 minute
	DHCPv4 Usage Trend DHCPv4 Range Utilization Trend	ib:dhcp:range	file-based (csv)	Overwritten every 1 hour
DHCP Lease History	DHCP Lease History DHCP Top Lease Clients	ib:dhcp:lease_history	file-based (syslog)	Rotates at 120 MB; retains one older copy; queued data is between 120 MB and 240 MB
	Top Devices Identified Device Trend Device Class Trend Top Device Classes	ib:dhcp:lease_history	file-based (syslog)	Based on summary search report, which is updated during the 16th and 46th minutes of each hour
	Top Devices Denied an IP Address	ib:dhcp:lease_history	file-based (syslog)	Based on summary search report, which is updated during the 19th and 49th minutes of each hour
	Device Fingerprint Change Detected	ib:dhcp:lease_history	file-based (syslog)	Executed every 24 hours

Report Category	Reports	Source Type	Data Source (file-based or scriptbased)	Update Frequency
DNS Performance	DNS Response Latency Trend	ib:dns:perf	script-based	Executed every 1 minute
DNS Record Scavenging	DNS Scavenged Object Count Trend	ib:dns:reclamation	file-based (csv)	Updated whenever reclamation tasks are executed
DNS Query Capture	DNS Domain Query Trend DNS Domains Queried by Client Top DNS Clients by Query Type Top DNS Clients Querying MX Records	ib:dns:capture	file-based (csv)	Updated whenever the Data Collection VM collects capture query data from a Grid member
DDNS	DDNS Update Rate Trend	ib:ddns	file-based (syslog)	Rotates at 120MB; retains one older copy; queued data is between 120MB and 240MB.
DNS Traffic Control	DNS Traffic Control Resource Availability Trend	ib:dns:reserved	file-based (csv)	Based on summary search report, which is updated once per six hour at 47th minute of each hour. With each execution, it summarizes raw events indexed from 370 minutes ago to 10 minutes ago.
	DNS Traffic Control Resource Availability Status	ib:dns:reserved	file-based (csv)	Based on summary search report, which is updated once per six hour at 47th minute of each hour. With each execution, it summarizes raw events indexed from 370 minutes ago to 10 minutes ago.
	DNS Traffic Control Resource Pool Availability Trend	ib:dns:reserved	file-based (csv)	Based on summary search report, which is updated once per six hour at 23rd minute of each hour. With each execution, it summarizes raw events indexed from 370 minutes ago to 10 minutes ago.
	DNS Traffic Control Resource Pool Availability Status	ib:dns:reserved	file-based (csv)	Based on summary search report, which is updated once per six hour at 23rd minute of each hour. With each execution, it summarizes raw events indexed from 370 minutes ago to 10 minutes ago.
	DNS Traffic Control Response Distribution Trend	ib:dns:reserved	file-based (csv)	Based on summary search report, which is updated once per six hour at 37th minute of each hour. With each execution, it summarizes raw events indexed from 370 minutes ago to 10 minutes ago.
DDI Utilization	DHCPv4 Usage Statistics DHCPv4 Top Utilized Networks	ib:dhcp:network	file-based (csv)	Overwritten every 1 hour
	IPAM Network Usage IPAM Top Networks	ib:ipam:network	file-based (csv)	Overwritten every 1 hour
	DNS Zone Statistics Per DNS View	ib:dns:view	file-based (csv)	Overwritten every 24 hours

Report Category	Reports	Source Type	Data Source (file-based or scriptbased)	Update Frequency
	DNS Statistics per Zone	ib:dns:zone	file-based (csv)	Overwritten every 24 hours
	DNS Object Count Trend for Flex Grid License	ib:dns:ibflex_zone_counts	file-based (csv)	Generated once in 24 hours and average is calculated over 5 days
System Utilization	CPU Utilization Trend Memory Utilization Trend Traffic Rate by Member	ib:system	script-based	Executed every 1 minute
	License Pool Utilization	ib:system	file-based (csv)	Overwritten every 24 hours
	SPLA Grid Licensing Features Enabled	ib:system		Generated once in 24 hours for all IB-FLEX members on the Grid
System Capacity	System Capacity Prediction	ib:system_capacity:objects		Updated whenever there is relevant event occurs
DNS Query	DNS Replies Trend	ib:dns:stats	script-based	Executed every 1 minute
	DNS Cache Hit Rate Trend	ib:dns:query:cache_hit_rate	script-based	Executed every 1 minute
	DNS Query Rate by Query Type	ib:dns:query:qps	script-based	Executed every 1 minute
	DNS Query Rate by Member DNS Daily Query Rate by Member DNS Daily Peak Hour Query Rate by Member	ib:dns:query:by_member	script-based	Executed every 1 minute
	DNS Top Clients	ib:dns:query:top_clients	script-based	Executed every 10 minutes
	DNS Top Requested Domain Names	ib:dns:query:top_requested_domain_names	script-based	Executed every 10 minutes

Report Category	Reports	Source Type	Data Source (file-based or scriptbased)	Update Frequency
	DNS Top Clients Per Domain DNS Top NXDOMAIN / NOERROR (no data) DNS Top SERVFAIL Errors Received DNS Top SERVFAIL Errors Sent DNS Top Timed-Out Recursive Queries	ib:dns:reserved	script-based	Executed every 10 minutes
	DNS Query Trend per IP Block Group	ib:dns:reserved	script-based	Executed every 5 minutes
	DNS Effective Peak Usage Trend for Flex Grid License	ib:dns:query:qps		Executed every 10 minutes and average is calculated over five days
Security	DNS Top RPZ Hits	ib:dns:reserved	script-based	Executed every 10 minutes
	DNS Top RPZ Hits by Clients	ib:dns:reserved	script-based	Executed every 10 minutes
	Top DNS Firewall Hits	ib:dns:reserved	script-based	Executed every 10 minutes
	Malicious Activity by Client	ib:dns:reserved	script-based	Executed every 10 minutes
	DNS Firewall Executive Threat	ib:dns:reserved	script-based	Executed every 10 minutes
	FireEye Alerts	ib:syslog	script-based	Updated immediately when alerts are logged in the syslog.
	Threat Protection Event Count By Severity Trend Threat Protection Event Count By Member Trend Threat Protection Event Count By Rule Threat Protection Event Count By Time Threat Protection Event Count By Category Threat Protection Event Count By Member	ib:reserved1	file-based (csv)	Overwritten every 5 minutes.

Report Category	Reports	Source Type	Data Source (file-based or scriptbased)	Update Frequency
	DNS Top Tunneling Activity DNS Tunneling Traffic by Category Top Malware and DNS Tunneling Events by Client	ib:reserved 1	file-based (csv)	Overwritten every 5 minutes.
Network User	User Login History	ib:reserved 1	file-based (csv)	
Ecosystem Subscription	Subscription Data	ib:reserved 1	file-based (csv)	Updated whenever there is an event received from the vendor that NIOS subscribes.
Ecosystem Publication	Publish Data	ib:reserved 1	file-based (csv)	Updated whenever there is a relevant RPZ, IPAM, and DHCP lease event occurs.
Cloud	VM Address History	ib:reserved 2	file-based (csv)	Updated immediately when there is a change related to the VM IP address. Rotates at 300MB and retains one older copy.
Audit Log	Audit Log Events	ib:audit	file-based (audit log)	Updated immediately when the audit log is updated.
	Audit Log WAPI Events	ib:audit	file-based (audit log)	Updated immediately when the audit log is updated.
Syslog	Syslog Events	ib:syslog	file-based (Syslog)	Updated immediately when alerts are logged in the syslog.

## Configuring Reporting Clusters

You can configure a reporting single indexer, a single-site cluster, or a multi-site cluster. When you configure reporting clustering, make sure that you configure two or more reporting appliances and that all indexers are online.



### Note

There is no action required if you see intermittent "Too many streaming errors" and "Skip indexing" messages in the **Messages** menu of the **Reporting** tab. This can be caused by network connectivity issues between the reporting nodes.

During NIOS upgrade, when configuring reporting clusters, ignore the "Unable to establish a connection to peer" message displayed on the Reporting tab.

To configure a reporting cluster:

1. From the **Administration** tab -> **Reporting** tab, click **Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **Reporting Clustering** tab and complete the following:
  - **Single Indexer**: Select this to configure only one reporting server. This is the default reporting cluster mode.

- **Single-Site Cluster:** Select this if you want to configure two or more reporting servers in the same site (location). The data is replicated on multiple reporting servers. You can upgrade your configuration to the multi-site clustering mode, but you cannot revert this configuration to a single indexer mode.
- **Multi-Site Cluster:** Select this if you want to configure multiple reporting servers at different sites (locations). You must assign the ReportingSite extensible attributes to all the reporting members that you have configured in the same site within the cluster. You can configure the same ReportingSite extensible attribute with multiple reporting members. The reporting members that are configured with the same ReportingSite extensible attributes are tagged to the same site. Click the Add icon and select the ReportingSite extensible attribute that you have configured on the reporting member. The first site that you add is considered to be the primary site, which functions as the search head. You can change the order of the sites by clicking the up and down arrows.

For more information about the reporting cluster type, see [Reporting Cluster Modes](#) above.

#### **Note**

Your multi-site configuration is invalid if you do not add the correct ReportingSite extensible attribute values to the reporting members. You can validate your configuration as described in [Validating Reporting Clustering Configuration](#) below.

3. Click **Save & Close**.

## Assigning a ReportingSite EA Value to a Multi-Site Cluster

The multi-site clustering configuration is valid only when you associate all the reporting members in the cluster with the specified **ReportingSite** extensible attribute values. Make sure that you select the ReportingSite values from those that are specified for the multi-site cluster in the *Grid Reporting Properties* editor. After you assign extensible attribute values to the reporting members, you can validate the multi-site cluster configuration as described in [Validating Reporting Clustering Configuration](#) below.

To associate the ReportingSite extensible attribute with the reporting member:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab -> *member* checkbox, and then click **Extensible Attributes** in the Toolbar.
2. Click the Add icon in the **Extensible Attributes** table to enter extensible attributes. The appliance adds a row to the table each time you click the Add icon. Select the row and the attribute name from the drop-down list, and then enter the value.
3. Optionally, select an extensible attribute and click **Delete** to delete it.
4. Click **Save & Close**.

## Validating Reporting Clustering Configuration

After you have configured the reporting cluster mode, you can verify its validity. Whenever you make changes to the reporting configuration through Grid Manager or hardware replacement, make sure that you validate the configuration. When you verify a multi-site cluster configuration, NIOS validates the following:

- The extensible attribute ReportingSite is specified for all reporting members.
- The set of extensible attributes configured in the *GridReportingProperties* editor equals to the set of ReportingSite extensible attributes defined for the reporting members.
- For each ReportingSite extensible attribute, the number of reporting peers must be greater or equal to the replication factor in each site.
- For each ReportingSite extensible attribute, the search factor must be less than or equal to the replication factor in each site.

To verify the reporting cluster-mode configuration:

1. From the Grid tab -> **Grid Manager** tab, click the **Reporting** service.
2. In the vertical Toolbar, click **Verify Cluster Configuration**.  
The *Verify Reporting Cluster Configuration* dialog box displays an error message if the configuration is invalid.

Make sure that you associate the ReportingSite extensible attributes with all the reporting members that you have configured.

3. Click **OK** to close the dialog box.

## Guidelines for Deploying Reporting Clusters

The following are a few guidelines to consider before you deploy reporting clustering:

- [Best Practices for Capacity Planning](#)
- [Monitor Volume and Storage Consumption](#)
- [Deployment and Service Monitoring](#)
- [Other Best Practices for Deploying Reporting Clusters](#)

### Best Practices for Capacity Planning

- Forecast the usage of the reporting volume and disk size based on your business needs and estimate the future potential growth.
  - Work with the your Infoblox representatives to gather statistics and metrics of your Grid to estimate the expected potential reporting data volume and disk storage required to manage your reporting data. This helps you determine which tier of NIOS Reporting license you should purchase, the number of appliances, and the appliance models you might need.
  - Configure reporting appliances using the same hardware appliance models. The capacity of the whole cluster is limited by the weakest reporting member if you configure reporting members using different appliance hardware models.
  - Consider the future potential growth during the initial deployment of a reporting cluster. Scaling out the reporting cluster by adding more reporting members at a later time might result in a significant configuration change.
  - Unlike single-indexer reporting solution, there might be more than one solution to handle certain estimated license and disk size for the reporting cluster. To reduce the potential admin overhead, Infoblox recommends that you use a small number of high-end reporting appliances to form a cluster, rather than using a large number of low-end reporting appliance.
- When deploying a reporting cluster in a virtual environment, consider the following:
  - You can deploy clustering in a virtual environment as long as it provides adequate system resources and capacities. Refer to the following resource to help design the clustering deployments. This link provides information about best practices and consequences that you must be aware of when configuring the host environments:
    - <https://docs.splunk.com/Documentation/Splunk/8.2.4/Capacity/Referencehardware>
- If you expect to deploy the reporting cluster in multiple data centers to achieve disaster recovery in the future, then configure the multi-site reporting clustering for the first time, even if all reporting members are physically located in the same data center. You can move some of the reporting members to other data center at a later time as required, rather than deploying new reporting members and reconfiguring the multi-site cluster.  
Note: Be aware that the additional reporting members (the secondary site in the same data center) doesn't really increase the capacity of the reporting cluster.
- Reporting cluster can only tolerate single node failure.
  - By default, the replication factor is set to 2, which means that there are two copies for each bucket in the entire single-site cluster or each site for a multi-site cluster. Therefore, it can tolerate failure only on one node.
  - The cluster may not be able to tolerate double failure. For example, if node1 just recovered from a failure and if node2 fails before node1 is fully recovered (before turning everything to green). In such cases, there is a chance of losing some data.
- When you change the reporting cluster mode, be aware of the following:
  - The data indexed by a single-indexer is not replicated to other peers, even if upgraded to a single-site cluster. Similarly, the data indexed by a single-site cluster is not replicated to other sites, even if upgraded to a multi-site cluster. In other words, the single-indexer configuration does not support HA configuration and single-site data does not have disaster recovery measure. Note that when you restore the reporting



data created in a single-site cluster to a multi-site cluster, the data is restored only on the primary site of the cluster and are not replicated on the secondary site.

## Monitor Volume and Storage Consumption

- Monitor the license volume usage on the reporting server regularly using the following dashboards:
  - Monitor the **Today's License Usage** (GB) and **License Usage Trend Per Member** panels available in the [Home Dashboards](#).
  - For more information about the the following, see [Internal Dashboards](#):
    - Reporting License Usage
    - Reporting Volume Usage Trend per Category
    - Reporting Volume Usage Trend per Member
- Configure alerts in the Infoblox Reporting and Analytics App to trigger actions when data volume reaches a specified limit.
- Monitor the disk usage: You can monitor the disk usage for each index in the *Grid Reporting Properties* editor and adjust the breakdown among indexes to make sure that the data is retained for the desired period of time.
- Take actions when the license volume or disk usage approaches the limit.
  - Choose to obtain licenses that can manage your data volume.
  - Expand the cluster with additional reporting members to increase the disk volume. If there is a surge of reporting data generation before you upgrade your license, you may reduce the inbound volume by turning off reporting on certain Grid members or certain reporting categories.
- Do not ignore license violation warnings, such as GUI yellow banner and SNMP traps.
  - When reporting license usage approaches the limit, NIOS displays the warning messages in the yellow banner, log messages, and SNMP traps (if configured).
  - When you receive five (5) violation notifications in a rolling period of 30 days, you cannot view reports or configure report related functions. To avoid reporting interruptions, promptly address any violation notifications by contacting Infoblox Technical Support.  
Note that the reporting server continues to process incoming data during the violation state. However, you cannot view any report or manage any report related functions until you fix the violation issue.

## Deployment and Service Monitoring

- A reporting cluster requires that the Grid Master and reporting members have the same IPv4 or IPv6 configurations.
  - Use either IPv4 or IPv6 networks for all members. There might be an impact on the performance if you use both IPv4 and IPv6 networks. There are a few guidelines to consider before you deploy reporting clustering. For more information about the guidelines, see [Guidelines for Deploying Reporting Clusters](#).
- Be aware of the time lag in displaying the actual state in the reporting service status.
  - NIOS reporting monitors all communication between Splunk cluster master (on the Grid Master) and reporting members, between Grid members (forwarders) and reporting members, and among reporting members.
  - Grid Manager displays the service status to indicate any network issues. However, the status is not real-time and there might be some time lag up to 5 minutes from the moment when the issues occur or recover.
- Regularly backup your reporting data even for a multi-site cluster deployment. Infoblox recommends that you perform reporting data backup before changing your configuration from a single indexer to a single-site cluster or multi-site cluster and from a single-site cluster to a multi-site cluster. For information about backing up reporting data, see [Managing Reporting Data](#).
- Use the **Reporting Clustering Status** dashboard to monitor the reporting cluster status.
  - If the reporting cluster is functioning properly, then the overall status is green. The status turns red if there is any network outage or reporting member goes offline, which indicates that the replication factor or search factor is not met. It might take some time for Grid Manager to change the status to green even after you fix the issue because the cluster needs to replicate buckets among peers.
- There can be duplicated bucket data in the reporting backup FTP server.

- In the cluster mode configuration, there is a chance of uploading multiple copies of the same bucket because reporting data backup takes data from multiple reporting members simultaneously. However, the reporting data restore process eliminates the duplicate copies of such buckets.
- It is a good practice to back up the Infoblox Reporting and Analytics App regularly, even though the appliance backs up data from the running search head to the Grid Master periodically. By doing this, it is easier to recover the data automatically by the new search head in case of a search head failure. For information about Backing Up and Restoring the Infoblox Reporting and Analytics App, see [Managing Reporting Data](#).

## Other Best Practices for Deploying Reporting Clusters

- When you plan a single site or multi-site reporting cluster, make sure that you join all the relevant reporting members to the cluster immediately. Any delays in adding the reporting members might cause some data loss because peers are unable to start indexing immediately. In addition, Infoblox recommends that you validate the reporting clustering configuration before you start the reporting service, as described in Validating Reporting Clustering Configuration, see [Configuring Reporting Clustering](#).
- When you need to bring down any reporting members for reasons such as maintenance and the member will be back online soon, you can use the `set reporting_cluster_maintenance_mode` CLI command to suppress the cluster master to initiate any bucket fix-up activities. For more information about this command, refer to the *InfobloxCLIGuide*. After completing the maintenance, use the same (`set reporting_cluster_maintenance_mode`) CLI command to turn off the reporting cluster maintenance mode.
- When you upgrade from a previous release to NIOS 7.3.200 or later, the Grid will automatically be upgraded to the single indexer mode. You can change the reporting clustering configuration after the upgrade.
- When you upgrade from NIOS 7.3.200 to a later NIOS release, the Grid will retain the cluster type configured in NIOS 7.3.200.
- During a scheduled upgrade, reporting members will restart and become unavailable to receive events from other Grid members. To avoid the scenario that all reporting members might go offline about the same time, Infoblox suggests that you put reporting members in different upgrade groups and schedule different upgrade times for these upgrade groups. By doing so, there will always be some reporting members available to receive data from other Grid members. Thus, you can avoid potential data loss during scheduled upgrades.
- When you upgrade from a previous release to NIOS 7.3.200 or later, reporting data generated before you enable reporting clustering is replicated on all reporting members. Therefore, in cases when a reporting appliance is down or offline, you may experience some data loss (although this data will be captured in the reporting backup file). Infoblox recommends that you back up your reporting data before an upgrade or enabling reporting clustering.
- You cannot revert the NIOS software version after you have configured a single-site or multi-site cluster from the single indexer mode. For example, if you upgrade the Grid from NIOS 7.3.0 to NIOS 7.3.200 or later with reporting service enabled (only the single indexer mode is supported until NIOS 7.3.200), the Grid will automatically be upgraded to a single indexer mode. If you then configure a single-site or multi-site cluster after the upgrade, you cannot revert to a single indexer in the upgraded version. To configure a single indexer, you must revert the NIOS software version to NIOS 7.3.0, which may cause some data loss because not all data stored on different reporting members in the previous reporting cluster is consolidated to the single indexer.
- When you plan to move a reporting appliance from one Grid to another, you must first enable the reporting service in the new Grid and then join the reporting appliance.
- Infoblox recommends that you do not reboot the reporting appliance or restart the reporting service when performing reporting backups or restores. For information about Backing Up Reporting Data and Restoring the Reporting Database, see [Managing Reporting Data](#).
- Infoblox recommends that you use appliance models with the same capacity as peers in a cluster. If you set up a reporting cluster using heterogeneous models, the appliance displays a warning message and limits the capacity of the entire cluster based on the peer that has the smallest capacity. For information about how to validate the reporting configuration, see [Configuring Reporting Clustering](#). The following table provides example configurations for the single-site and multi-site clustering.

### Clustering Configuration Examples

Clustering Types	Configuration Details	Validation
Single-site clustering (Replication Factor = Search Factor =2)	Forwarder volume = 2GB per day Licensing Master: 1 X 5G Primary Site: 2 X TR-1405 (5G)	Yes
	Forwarder volume: 5GB per day Licensing Master: 1 X 10G Primary Site: 2 X TR-2205	Yes
	Forwarder volume: 8GB per day Licensing Master: 2 X 10G Primary Site: 4 X TR-2205	Yes
	Forwarder volume: 8GB per day Licensing Master: 1 X 20G Primary Site: 2 X TR-4015	Yes
	Forwarder volume: 8GB per day Licensing Master: 1 X 10G Primary Site: 2 X TR-2205	No 8GB exceeds the effective capacity of 7GB.
	Forwarder volume: 5GB per day Licensing Master: 1 X 10G Primary Site: 4 X TR-2205	No A single 10G license makes this cluster underused.
Multi-site clustering (Replication Factor = 2 Search Factor =1)	Forwarder volume: 5GB per day Licensing Master: 1 X 10G Primary Site: 2 X TR-2205, Backup Site: 2 X TR-2205	Yes
	Forwarder volume: 5GB per day Licensing Master: 1 X 10G Primary Site: 2 X TR-2205, Backup Site: 1 X TR-2205	No Infoblox recommends that you configure at least two reporting members in each site (location).
	Forwarder volume: 8GB per day Licensing Master: 1 X 10G Primary Site: 4 X TR-2205, Backup Site: 2 X TR-2205	Yes
	Forwarder volume: 8GB per day Licensing Master: 1 X 10G Primary Site: 4 X TR-2205, Backup Site: 1 X TR-2205	No Because total volume 8GB from forwarders exceed the effective capacity of 7GB.

## Grid Reporting Properties

After you set up a dedicated reporting appliance in your Grid, you must configure the Grid reporting properties so you can communicate with the reporting appliance and retrieve report data through the Grid Master. In addition, you must select the correct report categories in order for the reporting server to generate the correct data in corresponding reports, as described in [Configuring Grid Reporting Properties](#) below.

By default, only superusers can configure the Grid reporting properties. When you enable the Grid reporting service, all members transmit data to the reporting server. You can disable data transmission from specific members to the reporting server. Before using the reporting service, you must configure the remote server to export the search results, as described in [Reporting \(Index\) Storage Space](#) below. Once you configure the reporting server and enable the reporting service on Grid members, you can view and manage reports through the **Reporting** tab of Grid Manager.



## Note

- When you reset the appliance using the `reset all` CLI command or reset the database using the `reset database` CLI command, reporting configurations are not preserved. If you reset the appliance, you must configure Grid reporting properties and remote server settings to use the reporting service.
- Expired cookies in a Splunk session are not removed in the Firefox browser by default. Expired cookies also cannot be reused. However, there is no impact on functionality.

Complete the following to set up your reporting solution:

1. Configure general reporting properties, including the selection of report categories, as described in [Configuring Grid Reporting Properties](#) below.
2. Specify the network port for reporting, as described in [Setting the Network Port for Reporting](#) below.
3. Specify email properties, as described in [Configuring Email Notification Settings](#), [About Alerts](#).
4. Configure the logo image for PDF delivery, as described in [Configuring Logo Image in PDF Reports](#), see [About Reports](#).

The properties you define in the *Grid Reporting Properties* editor apply to all the reporting members unless you override them at the specific member level. To override at the member level, see [Modifying Member Reporting Properties](#) below.

## Configuring Grid Reporting Properties

After you configure the reporting server, you must enable the data indexing and select at least one reporting category to ensure that the reporting service functions properly.



## Note

You must select the correct report categories in order for the reporting server to generate the correct data in corresponding reports.

Complete the following to configure the Grid reporting properties:

1. From the **Administration** tab -> **Reporting** tab, click **Grid Reporting Properties** from the Toolbar.  
or  
From the **Grid** tab, select the **Grid Manager** tab and click the **Services** tab. In the **Services** tab, select the **Reporting** tab and click **Edit** -> **Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **General** -> **Basic** tab
3. Complete the following:
  - a. **Reporting Server**: Grid Manager displays the name of the reporting server.
  - b. **Enable Data Indexing**: Data transmission is disabled by default. You must select this checkbox to ensure that all Grid members transmit data to the reporting appliance. Enabling data transmissions for all members can affect the overall data consumption on the reporting server. For information about the daily maximum data consumption per day for your reporting appliance, see the [Default Index Space Configured for Each Report Category](#) table below.
  - c. **Enable Time Based Retention**: Select this checkbox if you want the reporting data to be retained for a number of days you specify. You can specify the number of days in the **Retention in days** column. When you select this checkbox, NIOS displays a warning message indicating that the reporting data is deleted after the number of days specified in the **Retention in days** column. This checkbox is disabled by default. Note that the **Enable Time Based Retention** feature is not available for the IP Address Usage Report, the DNS QPS Usage Report, and the DHCP LPS Usage Report as the deletion of data after the retention period impacts the subscription count, which in turn generates inaccurate data in these reports.

- d. **Report Category:** Select the reports you want the reporting server to generate. The reporting server automatically configures data sources and configurations required to generate the reports you select here. The required data is stored in the reporting server database. By default, no report categories are selected. For a list of report categories, see [Predefined Dashboards](#). You must select at least one reporting category for the reporting service to start working.
  - i. **Index%:** Displays the actual storage space allocated for a reporting index. You can modify this value between 0 and 100. When you enable an index category and leave it at 0%, the appliance displays an error message. Make sure that the total percentage of the index storage space for all report categories equals 100% or less than 100%. The appliance displays a warning message when the total percentage of the index storage space is less than 100%.
  - ii. **Used%:** Displays the index storage space used by a reporting index.
  - iii. **Retention in days:** Enter the number of days up to which you want the reporting data to be retained. The data will be permanently deleted after this number of days. You can enter a value between 7 and 365. The default value is **No Retention**.
  - iv. **IndexName:** Displays the reporting index name, which is displayed on the [Reporting Index Usage Statistics](#) report.  
Note that if you back up the reporting data before the expiry of the retention period, then you can restore the data at a later date.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Caching Threat Category Information from the Cloud Services Portal

The threat category information (Threat indicator database and Threat description) is downloaded from the Cloud Services Portal and stored locally. The threat category information is then sent to the reporting server to augment RPZ hits and reports are generated. Caching threat category information from the Cloud Services Portal helps enhance the performance of threat reports as data is fetched from the cache that is stored locally.

You can configure the Cloud Services Portal credentials and schedule the entire threat indicator database download from the Cloud Services Portal. If you have already downloaded the entire threat indicator database, then consecutive full downloads take place only after 24 hours.



### Note

For the threat indicator caching feature to work on a Grid, the Grid must have at least one user with **can delete** permission set up on the Grid.

## Limitations

Note the following limitations when you use the threat indicator caching feature:

- Enabling the threat indicator caching feature results in higher usage of network bandwidth and reduction of the reporting indexing capacity.
- Enabling the threat indicator caching feature impacts the performance of Grid Master as Splunk consumes significant bandwidth to forward the entries to indexers. It takes a few minutes for the entries to get forwarded and indexed completely based on the data size.
- If you enable the threat indicator caching feature, and then revert or upgrade the Grid to a version that does not support the feature, then the indexed threat indicator database data will still occupy disk space even though they are not searchable in the upgraded or reverted Grid version.
- The size of the downloaded threat indicator database file will be huge due to data retention in the following scenarios:
  - When you enable and disable the threat indicator caching feature a few times.
  - When you upgrade NIOS and then revert it to the prior version without disabling the threat indicator feature, and also when you upgrade NIOS again.
- When the threat indicator caching feature is enabled, threat details in the *DNS Top RPZ Hits* report does not show historic data. For more information about the *DNS Top RPZ Hits* report, see *DNS Top RPZ Hits*.

- For replication to work properly in cluster mode, Infoblox recommends that an appliance should have 12 cores CPU and 12 GB memory.

## Configuring the Threat Indicator Caching Feature

Complete the following to configure the threat indicator caching feature:

1. From the **Grid** tab, select the **Grid Manager** tab and click the **Services** tab. In the **Services** tab, select the **Reporting** tab and click **Edit** -> **Grid Reporting Properties** from the Toolbar.
2. In the **Grid Reporting Properties** editor, select the **Threat Indicator Caching** tab-> **Basic** tab.
  - a. **Enable Threat Indicator Caching:** Select the checkbox to enable the feature. Enabling this feature downloads the threat indicator information from the Cloud Services Portal to the Grid Master, and then the threat indicator information is indexed on the reporting members.  
Note that selecting this option results in higher indexing license usage, network bandwidth, and storage.
3. Complete the following:
  - a. **Splunk Threat Indicator Caching Index Storage:** Specifies the disk storage allocation for the threat indicator caching feature. The minimum disk storage limit is 8 GB and the maximum disk space that it can be set to is 42 GB. By default, the disk storage space is set to 12 GB. The disk space that you allocate will reduce the storage limit for all other indexes. Set the required storage space based on the volume of data that you expect to be downloaded from the Cloud Services Portal and based on your indexing capacity. Grid Master downloads the threat indicator data and periodically forwards it to the reporting server for indexing.  
The indexing usage that is observed by Infoblox during the lab testing is, one full synchronization consumes ~600 to ~800 MB of indexing space and each incremental synchronization consumes ~60 MB of indexing space.  
Note that the indexing space usage varies on a daily basis based on data generated by the Cloud Services Portal. Therefore, do not consider the numbers stated here as standard guidelines.
  - b. **Incremental Threat Indicator Caching Update Interval** (in hours): Enter the interval value in hours to download the incremental updates from the threat indicators of the Cloud Services Portal. For example, if you set the value as 2, after every two hours the incremental threat indicator is downloaded. The incremental threat indicator is downloaded only after the whole threat indicator is downloaded from the Cloud Services Portal.
  - c. **Last Incremental Threat Indicator Caching Download Timestamp:** Displays the date and time of the last successful incremental threat database download.
  - d. **Update Policy:** Select **Automatic** or **Manual**. You need to select any one of the following options in order to avoid huge data storage usage on Splunk.
    - i. **Automatic:** Select this option if you want to automatically download the whole database after every seven days. By default, the value is set to seven days.
    - ii. **Manual:** Select this option to schedule the whole database download manually. For more information on threat context locale cache scheduler, see [Scheduling Threat Indicator Caching](#) below.
    - iii. **Test Connection:** Click **Test Connection** to test the connectivity between NIOS and the Cloud Services Portal. Then, enter the Cloud Services Portal credentials on the **BloxOne Threat Defense Cloud Integration** tab. For more information about Enabling the BloxOne Threat Defense Cloud Client, see [Configuring BloxOne Threat Defense Cloud Clients for Outbound](#).
  - e. **Last Whole Threat Indicator Caching Download Timestamp:** Displays the date and time of the last successful whole threat indicator download.
  - f. **Scheduling:** Select to schedule the whole threat indicator download. You can select **Scheduling** only if the **Update Policy** is selected as **Manual**.
  - g. **Last Threat Indicator Caching Failure Timestamp:** Displays the date and time of the last failed attempt that is made to download the threat indicators after five iterations.

## Scheduling Threat Indicator Caching

You can schedule the download of the whole threat database daily, weekly or monthly. However, if you have already downloaded the whole threat indicator database and the scheduled date and time is near next, in that case the schedule is skipped. Based on the schedule the Incremental ThreatDB is downloaded from the Cloud Services Portal as per the set interval

Complete the following to schedule the threat indicator caching:

1. From the **Grid** tab, select the **Grid Manager** tab and click the **Services** tab. In the **Services** tab, select the **Reporting** tab and click **Edit** -> **Grid Reporting Properties** from the Toolbar.
2. In the **Grid Reporting Properties** editor, select the **Threat Indicator Caching**-> **Basic** tab
3. Select the **Enable Threat Indicator Caching** option.
4. In the *Complete ThreatDB Download Interval* section, select the **Update Policy** as **Manual**.
5. Click the Scheduling icon.
6. In the *Threat Indicator Caching Scheduler* dialog box, select the daily, weekly, or monthly option as follows:  
Note that the time zone for scheduling will be the same as the zone that is set for the Grid.
  - a. To schedule a daily download:
    - i. Select **Daily**.
    - ii. In the *Schedule daily* section, in the **Time** field, set the time when the download must start.  
Ensure that you have sufficient space available as daily data indexing limits may get exhausted.
  - b. To schedule a weekly download:
    - i. Select **Weekly**.
    - ii. In the *Schedule every week on* section select the day of the week when the download must happen every week.
    - iii. In the **Time** field, set the time when the download must start.
  - c. To schedule monthly download:
    - i. Select **Monthly**.
    - ii. In the *Schedule the day of the month* section, enter the day of the month when the download must happen. You can set the date from 1st to 28th of a month.
    - iii. In the **Time** field, set the time when the download must start.
  - d. Click **OK**.
7. Click **Save & Close**.

## Reporting (Index) Storage Space

One key configuration aspect of the reporting appliance is index space. By default, some percentage of index space is allocated on the reporting server for each report category listed in [Report Categories, Related Data Sources, and Update Frequencies table](#). For information about how to configure index space, see [Configuring Grid Reporting Properties](#) above. Each report category uses up to a certain percentage of the usable reporting hard disk space for index storage. For example, of the total 237 GB usable hard disk space of an IB-VM-800 appliance, the reporting category, **Device** uses 47.47%. For the list of default index space configured for each report category, see [Default Index Space Configured for Each Report Category table](#). You can modify the index percentage value between 0 and 100. When you modify this value, make sure that the total percentage for the index storage space for all categories equals exactly 100%. You can set the index percentage to a value of less than 100% to reserve a certain percentage for future use. If the total percentage of the index space usage exceeds 100%, the appliance displays an error message. Note that the reporting appliance removes the oldest data when you reduce the index space percentage for a category to a value that is lower than the used percentage by the existing data. For information about the maximum index size and number of days the reporting data is retained, see [Reporting Indexes table](#). Also, ensure that its host name has only alphanumeric characters, underscores, dots, and dashes.



### Note

For usable reporting hard disk space for each appliance model, see [Infoblox Reporting Appliances and their Usable Reporting Hard Disk Space table](#).

### *Default Index Space Configured for Each Report Category*

Report Category	Default Index Space (%) Adjustable by User	Total Reporting Disk Space Used for Index Storage (GB)
Audit Log	0%	-
DNS Query DNS Performance DDNS DNS Record Scavenging	20%	Usable reporting hard disk space x 20%
DNS Query Capture	0%	-
DHCP Performance	20%	Usable reporting hard disk space x 20%
DHCP Fingerprint DHCP Lease History	39%	Usable reporting hard disk space x 39%
DDI Utilization	5%	Usable reporting hard disk space x 5%
Security Network User	1%	Usable reporting hard disk space x 1%
DNS Traffic Control	0%	Usable reporting hard disk space is broken down between ib_dtc and ib_dtc_summary internally.
Cloud	0%	-
System Utilization	15%	Usable reporting hard disk space x 15%
		-
Device	0%	-
Ecosystem Subscription Ecosystem Publication	0%	-
License	0%	-

## Modifying Member Reporting Properties

To modify reporting properties for a reporting member:

1. From the **Grid** tab -> **Grid Manager** tab -> **Services** tab, select the **Reporting** service and click the *Grid\_member* checkbox, and then click the Edit icon.



2. In the *Reporting Member Properties* editor, select the **General** tab and click **Override**.
3. Under **Reporting Settings**, complete the following:
  - a. **Enable data forwarding to the indexer on this member:** Select this checkbox to enable data transmissions to the reporting server. If you do not select this checkbox, a member will not forward data to the indexer and reporting service is disabled on that member.
  - b. **Select the data categories to forward:** Select the report categories for which you want this member to forward data to the reporting server. Clear the report categories for which you do not want this member to forward data to the reporting server.  
Note that the member configured as an indexer displays only the **Audit Log** category.
4. Save the configuration.

## Defining Interface for Reporting Traffic

On a Grid member, you can define the network interface you want the member to use for sending reporting data to the reporting server.

To define network interface on the Grid member for reporting traffic, complete the following:

1. From the **Grid** tab -> **Grid Manager** tab -> **Services** tab, select the **Reporting** service and click the *Grid\_member* checkbox, and then click the Edit icon.
2. In the *Reporting Member Properties* editor, select the **General** -> **Advanced** tab, and then complete the following:
  - a. **Forwarding interface used for reporting traffic:** From the drop-down list, select the interface that you want this member to use to send reporting data. Note that you must properly configure the interfaces on the member for them to appear in the drop-down list. After a NIOS upgrade to version 8.1.x or later, if you had selected Any from the drop-down list, the LAN1 (or VIP for HA configurations) subnet is used as the static route except if the threat protection licenses are installed and ANY is selected from the drop-down, MGMT subnet is used as the static route. Select MGMT if you want to continue using the management port for grid communication. Selecting MGMT enables the MGMT subnet to be used as the static route.  
Note that after you start the reporting service on the reporting member, you cannot reset the interface set for the reporting traffic. You may have to configure the reporting member again to modify the interface for the reporting traffic.
3. Save the configuration.

## Setting the Network Port for Reporting

All Grid members use port 9997 for reporting service by default. This port is used for data transmissions between the reporting member and other members. Ensure that you configure your firewall rules to allow traffic on this port. You can designate another network port for reporting purposes.

To set the network port for reporting, complete the following:

1. From the **Administration** tab, select the **Reporting** tab -> expand the Toolbar and click **Grid Reporting Properties**.  
Or  
From the **Grid** tab -> **Grid Manager** tab -> **Services** tab, select the **Reporting** service and click the *Grid\_member* checkbox, and then click **Edit** -> **Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **General** -> **Advanced** tab and complete the following:
  - a. **Port:** Enter the port number you want to use for reporting purposes. The default port is 9997.
3. Save the configuration.

## Specifying the Data Generation Interval for Reports

You can specify the time interval when the appliance generates data for the *DNS Statistics per View* and *DNS Statistics per Zone* reports. The default value for the data generation interval for these reports is one day (86400 seconds). You can specify different data generation intervals for the *DNS Statistics per View* and *DNS Statistics per Zone* reports.

To specify the data generation interval for *DNS Statistics per View* and *DNS Statistics per Zone* reports, complete the following:

1. From the **Administration** tab, select the **Reporting** tab -> expand the Toolbar and click **Grid Reporting Properties**.  
or  
From the **Grid** tab -> **Grid Manager** tab and click the **Services** tab. In the **Services** tab, select the **Reporting** tab and click **Edit -> Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **Data Generation Schedule** tab and complete the following:
  - a. **Data Generation**: Enter the time in HH:MM:SS AM/PM format. You can also click the Clock icon to select a time from the drop-down list.
3. Save the configuration.

For more information about the reports *DNS Statistics per DNS View* and *DNS Statistics per Zone*, see [IPAMv4 Utilization Dashboards](#).

## Configuring Threat Protection Data

You can use this feature only if the Threat Protection and Threat Protection Update licenses are installed on the Infoblox Advanced Appliance. When you configure this feature, you receive threat protection events in the syslog. The events logged include threat protection rules and the source IPs that triggered the rules. For information about how to monitor these events using the syslog, see [Monitoring through Syslog](#).

For certain threat protection reports, accumulated statistics for each unique IP/rule pair are collected. You can control the volume of data collected per member using the following options:

- **Top IP/Rule Statistics Collection Limit**: This option limits the collection of accumulated statistics to the top N unique IP/rule pairs.
- **IP/Rule Statistics Collection Interval (minutes)**: The interval at which the accumulated statistics for the top N unique IP/rule pairs are collected. The smaller the interval, the finer the granularity of the accumulated statistics in terms of time, but the data volume will be higher.

Based on your configuration, the reporting appliance displays data in the following threat protection reports, see [Security Dashboards](#):

- Threat Protection Top Rules Logged
- Threat Protection Top Rules Logged by Source



### Note

When threat details are missing for a non-local RPZ feed zone entry, it is recommended to check if the associated feed zone's TSIG key is configured.

To enable threat protection reports, you must select the **Security** report category in the **Grid Reporting Properties** editor. To select the **Security** checkbox, go to the **Reporting** tab -> **Grid Reporting Properties** -> **General** tab -> **Basic** tab -> select the **Security** checkbox under **Report Category**. Ensure that you set the Security Index% to an optimal level so the reporting database has enough storage space to accommodate all reporting data. For information about how to configure the Index%, see [Configuring Grid Reporting Properties](#).

To configure the data collection limit, complete:

1. From the **Administration** tab, select the **Reporting** tab and click **Grid Reporting Properties** from the Toolbar.  
Or  
From the **Grid** tab, select the **Grid Manager** tab and click the **Services** tab. In the **Services** tab, select the **Reporting** tab and click **Edit -> Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **Threat Protection** tab and complete the following:
  - a. **Top IP/Rule Statistics Collection Limit**: Enter the maximum number of the top N unique source IP/rule pairs for data collection. For example, if you specify 20, the appliance collects data for the top 20 unique source IP/rule pairs.

- b. **IP/Rule Statistics Collection Interval (minutes)**: Enter the time interval at which the reporting appliance updates data. For example, if you specify the interval as 60 minutes, the appliance updates data at a 60-minute interval.
3. Click **Save & Close**.

## Monitoring DNS Client Queries

You can view the presence of clients in the network that are sending large numbers of queries to DNS zones or DNS domains. To monitor the top clients querying DNS zones, perform the following:

1. From the **Administration** tab, select the **Reporting** tab -> expand the Toolbar and click **Grid Reporting Properties**.  
Or  
From the **Grid** tab, select the **Grid Manager** tab and click the **Services** tab. In the **Services** tab, select the **Reporting** tab and click Edit -> **Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **Basic** tab -> **DNS**.
3. Under **DNS Top Clients Per Domain**, select the **Monitor Queries made to the following zones** checkbox. Only authoritative zones are supported, to a limit of 1000 zones for monitoring purposes.
  - a. To select zones one at a time, choose individual checkboxes. Click the Add icon and select **Add Domain** or **Bulk Add Domains** to add new zone information for excluding.
  - b. To specify the number of clients to be listed, choose the **Top N Limit** value. The default value is 10.

## Monitoring IP Block Group Queries

You can view the user defined IP block groups that are querying DNS domains. To monitor the IP Block Groups, perform the following:

1. From the **Administration** tab, select the **Reporting** tab -> expand the Toolbar and click **Grid Reporting Properties**.  
Or  
From the **Grid** tab, select the **Grid Manager** tab and click the **Services** tab. In the **Services** tab, select the **Reporting** tab and click **Edit** -> **Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **Basic** tab -> **DNS**.
3. Under **DNS Query trend per IP Block**, select the **Monitor Queries made from the following groups** checkbox.
  - a. Click the Add icon to add a group to the group table. From the drop-down list, click **Select Group** to select groups in the *Group Selector* dialog box, or click **Bulk Add Groups** to add multiple groups.
  - b. To select all groups, select the **Group** checkbox. Or, select individual checkbox to select the group one at a time.
  - c. To delete a group, select the group and click the Delete icon.

## Configuring DNS RPZ Rule Hits

You can specify a limit to display the number of top clients, who receive re-written responses through the RPZ, in **DNS Top RPZ Hits**. You can also specify the total number of RPZ entries for each client.

1. From the **Administration** tab, select the **Reporting** tab -> expand the Toolbar and click **Grid Reporting Properties**.  
Or  
From the **Grid** tab, select the **Grid Manager** tab and click the **Services** tab. In the **Services** tab, select the **Reporting** tab and click **Edit** -> **Grid Reporting Properties** from the Toolbar.
2. In the *Grid Reporting Properties* editor, select the **Basic** tab -> **DNS**.
3. Under **DNS RPZ Rule Hit Configuration**, complete the following:
  - a. Enter a value for **Top N Limit** to specify the maximum number of top clients that can be listed in the report.
  - b. Specify the **Total RPZ Entries per Client**. This indicates the number of entries for each client in RPZ.



### Note

You have to select the **Security** checkbox before you define values here. To select the checkbox, **Reporting** tab -> **Grid Reporting Properties** -> **General** tab -> **Basic** tab -> select the checkbox **Security** under **Report Category**.

## Forwarding Syslog Data to the Reporting Server

You can control the kind of syslog data forwarded to the indexer from the Grid members. You can search for syslog events (search string) in the **Reporting** tab -> **Search** tab. The syslog events you see in the **Search** tab depend on the syslog categories that you specify in both the Grid Reporting and Member Reporting Properties. The **Search** tab displays syslog events for the selected syslog categories at both the Grid Reporting Properties and Member Reporting Properties.

To specify syslog data categories, complete the following:

1. From the **Administration** tab, select the **Reporting** tab -> expand the Toolbar and click **Grid Reporting Properties**.  
Or  
**Member**: From the **Grid** tab, select the **Grid Manager** tab and click the **Services** tab. In the **Services** tab, select the **Reporting** tab -> *member* checkbox and then click the Edit icon.
2. In the *Grid Reporting Properties* or *Reporting Member Properties* editor, select the **Syslog Data** tab and complete the following:  
Click **Override** in the *Reporting Member Properties* editor to override the settings configured at the Grid reporting level. To inherit the same properties as the Grid, click **Inherit**.
  - a. **Source**: From the drop-down list, select which syslog messages the appliance sends to the external syslog server:
    - i. **Any**: The appliance sends both internal and external syslog messages.
    - ii. **Internal**: The appliance sends syslog messages that it generates.
    - iii. **External**: The appliance sends syslog messages that it receives from other devices, such as syslog servers and routers.
  - b. **Severity**: Choose a severity filter from the drop-down list. When you choose a severity level, the appliance sends log messages with the selected level and the levels above it. The severity levels range from the lowest, **debug**, to the highest, **emerg**. For example, if you choose **debug**, the appliance sends all syslog messages to the server. If you choose **err**, the appliance sends messages with severity levels **err**, **crit**, **alert**, and **emerg**.
    - i. **emerg**: Panic or emergency conditions. The system may be unusable.
    - ii. **alert**: Alerts, such as NTP service failures, that require immediate actions.
    - iii. **crit**: Critical conditions, such as hardware failures.
    - iv. **err**: Error messages, such as client update failures and duplicate leases.
    - v. **warning**: Warning messages, such as missing keepalive options in a server configuration.
    - vi. **notice**: Informational messages regarding routine system events, such as "starting BIND".
    - vii. **info**: Informational messages, such as DHCPACK messages and discovery status.
    - viii. **debug**: Messages that contain information for debugging purposes, such as changes in the latency timer settings and AD authentication failures for specific users.
  - c. **Logging Category**: Select one of the following logging categories:
    - i. **Send all**: Select this to log all syslog messages, irrespective of categories to which it belongs. When you select this option, the appliance logs syslog messages for all the events, including all DNS and Infoblox related events. However, the syslog messages are not prefixed when you select this option.
    - ii. **Send selected categories**: Select this to configure logging categories from the list of available logging categories. Use the arrows to move logging categories from the **Available** table to the **Selected** table and vice versa. The appliance sends syslog messages for the categories that are in the **Selected** table. When you select this option, you must add at least one logging category. The syslog messages are prefixed with a category name to which it belongs. Also, the RPZ events logged in the syslog messages uses specific prefixes for the selected categories. Note that the syslog messages are prefixed when you set logging categories for at least one external syslog server, even if you set other external syslog servers as **Send All**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring IP Blocks and IP Block Groups

You can configure IP addresses, subnets, or a mix of multiple IP addresses and subnets into IP blocks, and then assign them to IP block groups for monitoring and tracking queries made to specific IP blocks. For information about adding IP addresses, see [Adding IP Blocks](#) below. You can also configure as many groups as necessary and assign them to specific clients. Note that assigning more IP block groups results in monitoring more queries, which may affect the performance of the reporting server. You can generate a report to monitor queries made to these user-defined IP block groups or IP blocks. For information about DNS Query Trend per IP Block Group, see [Predefined Dashboards](#). Guidelines while configuring IP blocks:

- You cannot configure arbitrary IP address ranges, such as 192.168.0.1 to 192.168.0.100 as an IP block.
- You cannot add or modify an IP block that overlaps with another IP block in a different group. However, you can add an IP block that overlaps with another IP block in the same IP block group.



### Note

The appliance restarts the DNS service after you assign or unassign an IP block group at the Grid or member level. Also, the appliance restarts the DNS service when you modify or delete an IP address or IP block group assigned to the Grid or member or when you add, modify, or delete an IP block in such IP block groups.

You can do the following in the Groups panel, as described in the following sections:

- Add IP block groups.
- Modify IP block groups.
- Add IP block.
- Modify IP blocks.
- Delete IP block groups and IP blocks.
- Print IP block groups and IP blocks.
- Export IP block groups and IP blocks.

In addition, you can also do the following:

- Use filters or the **Go to** function to navigate to a specific group. You can also create quick filters to save frequently used filter criteria. For information about how to use quick filters, see [Finding and Restoring Data](#).
- Use Global Search to search for IP block groups and IP blocks. For information about Global Search, see [About the Grid Manager Interface](#).
- Use Smart Folders to organize IP block groups and IP blocks. For information, see [Smart Folders](#).
- Import and export groups in CSV format. For more information about CSV import feature, see [Importing and Exporting Data using CSV Import](#).

## Adding IP Block Groups

To add a new group:

1. From the **Administration** tab, select the **Reporting** tab -> *Group* -> Add.  
or  
From the **Administration** tab, select the **Reporting** tab, expand the Toolbar and click **Add** -> IP Block Group.
2. In the *Add IP Block Group* wizard, complete the following:
  - **Name**: Enter the name of the group.
  - **Comment**: Enter useful information about the group.
3. Do one of the following:
  - Click **Save & Close** to add the IP block group and close the wizard.
  - Click **Save & Edit** to add the IP block group and launch the editor. You can edit the details.
  - Click **Save & New** to add the IP block group and launch the wizard again to add another group.
  - Click **Save & Open** to add the IP block group and open the IP block group.

## Modifying IP Block Groups

1. From the **Administration** tab, select the **Reporting** tab -> *Group* -> Edit.
2. In the **General** tab of the *IP Block Group* editor, you can modify the group name and comment.
3. Click **Save & Close** to save the configuration.

### **Note**

You can perform inline editing by double-clicking the row of data that you want to modify. The appliance displays the inline editing editor in the selected row. Click **Save** after modifying the data.

## Adding IP Blocks

In a group, you can add as many subnets/IP addresses as necessary. Note that adding more IP addresses results in monitoring more queries, which may affect the performance of the reporting server.

1. From the **Administration** tab, select the **Reporting** tab -> *Group* -> Add.  
or  
From the **Administration** tab, select the **Reporting** tab, expand the Toolbar and click **Add** -> IP Block.
2. In the *Add IP Block* wizard, complete the following:
  - **Group:** Click **Select** to select a group. When there are multiple groups, Grid Manager displays the *Group Selector* dialog box to select a group. Click a group name in the dialog box. You can use filters or the Go to function to narrow down the list.
  - **Address:** Enter the source IPv4/IPv6 addresses or the IPv4/IPv6 subnets.
  - **Comment:** Enter useful information about the IP block.
3. Do one of the following:
  - Click **Save & Edit** to add an IP address or IP block and launch the editor. You can edit the details.
  - Click **Save & New** to add an IP address or IP block and launch the wizard again to add another IP block.
  - Click **Save & Close** to add an IP address or IP block and close the wizard.

## Modifying IP Blocks

1. From the **Administration** tab, select the **Reporting** tab -> *Group*.
2. Select an IP address or IP block you want to modify and click the Edit icon.
3. In the **General** tab of the *IP Block* editor, modify the IP address or comment.
4. Click **Save** to save the configuration.

### **Note**

You can modify description by using inline editing. Double-click the row that you want to modify, the appliance displays the inline editing editor in the selected row. Click **Save** after modifying comment. You cannot modify IP address using inline editing editor.

## Deleting IP Block Groups and IP Blocks

1. For IP block groups: From the **Administration** tab, select the **Reporting** tab -> *Group*.
2. For IP blocks: From the **Administration** tab, select the **Reporting** tab -> *Group* -> *IP block*.
3. Click the Delete icon.
4. In the *Delete Confirmation* dialog box, click **Yes**.

## Exporting IP Block Groups and IP Blocks

You can export displayed data or you can export the group list in CSV (comma separated value) format. Exporting group lists or group data may take a few moments based on the amount of exported data.

To export displayed data:

1. For IP block groups: From the **Administration** tab, select the **Reporting** tab -> *Group*.
2. For IP blocks: From the **Administration** tab, select the **Reporting** tab -> *Group -> IP block*.
3. From the **Export** drop-down menu, select **Export visible data**. For more information on how to export, see [Importing and Exporting Data using CSV Import](#).

To export all data to a CSV file:

1. For IP block groups: From the **Administration** tab, select the **Reporting** tab -> *Group*.  
For IP blocks: From the **Administration** tab, select the **Reporting** tab -> *Group -> IP block*.
2. From the Export drop-down menu, select **Export data in Infoblox CSV Import format**. For more information on how to export, see [Importing and Exporting Data using CSV Import](#).

## Printing IP Block Groups and IP Blocks

1. For IP block groups: From the **Administration** tab, select the **Reporting** tab -> *Group*. For IP blocks: From the **Administration** tab, select the **Reporting** tab -> *Group -> IP block*.
2. Click the Print icon. For more information on how to print from Grid Manager, see [About the Grid Manager Interface](#).

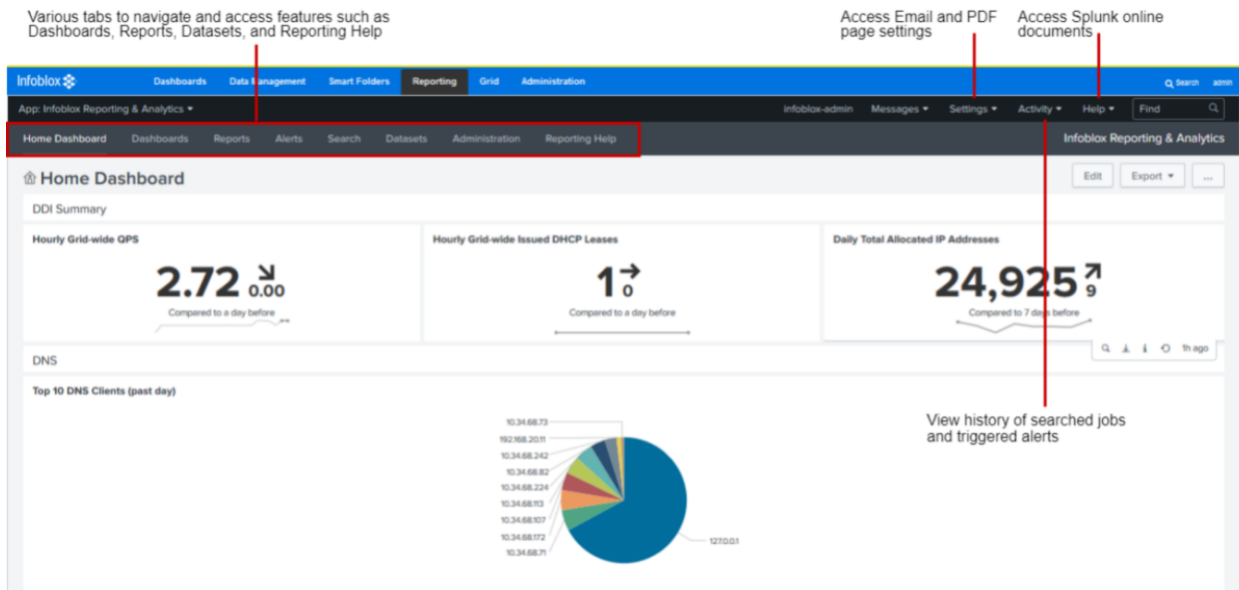
## Reporting User Interface Overview

When you access the **Reporting** tab in Grid Manager, you can do the following:

- View and analyze all of your DNS, DHCP, and IPAM data.
- View the reporting data in the form of charts and tables.
- Use the **Search** tab to create or edit searches.

The key options in the **Reporting** tab are: **Home Dashboards**, **Dashboards**, **Reports**, **Alerts**, **Search**, **Datasets**, **Administration**, and **Reporting Help**. For more information about these options, refer the below table.

*Reporting Tab Overview*



The following table summarizes the different tabs and key options that are available on the **Reporting** tab:  
*Reporting Tab Overview*

Tab Name	Purpose	Description
<b>Settings</b>	Email notification settings Server Logging	For information about Configuring Email Notification Settings, see <a href="#">About Alerts</a>
<b>Home Dashboard</b>	Overall summary view for the reporting data in your Grid	For information, see <a href="#">Home Dashboards</a>
<b>Dashboard</b>	Provide table and graph views for the reporting data in your Grid.	For information, see <a href="#">About Dashboards</a>
<b>Reports</b>	Saved searches that retrieve specific type of reporting data	For information, see <a href="#">About Reports</a>
<b>Alerts</b>	Set alerts to trigger actions when certain events occur	For information, see <a href="#">About Alerts</a>
<b>Search</b>	Create a search interactively from scratch and save it as a dashboard panel, alert and report.	For information about About Searches, see <a href="#">Home Dashboards</a>
<b>Datasets</b>	Refer to Splunk documentation for more information.	
<b>Administration</b>	<ul style="list-style-type: none"> <li>Set up</li> <li>Permissions</li> </ul>	For information, see <a href="#">Reporting (Index) Storage Space</a>  For information, see <a href="#">Administrative Permissions</a>
<b>Reporting Help</b>	Quick help on some of the reporting topics.	



Hierarchy settings for all the navigation menu options (listed in the Reporting Tab Overview table above) are available in the *default.xml* file.

 **Warning**

*Infoblox recommends that you do not modify the default.xml file. The navigation menu is built on a custom XML structure that is stored as default.xml in the navigation directory (from the **Reporting** tag, select the **Settings** tab -> **User Interface** -> **Navigation Menus** -> default). Editing this file changes the specific default settings for the reporting features and your changes become permanent. In addition, you might not be able to see the latest changes made by Infoblox. To restore the default settings, you must reset the NIOS appliance to its original factory settings. For Information, see [Resetting a NIOS Appliance to Factory Settings](#).*

## Predefined Dashboards

Grid Manager includes a list of predefined dashboards that provide summary views for most of the data and trends in your Grid. If you need to modify the settings of a default dashboard, you can either clone the default dashboard, or create a new one from scratch to create a user-defined dashboard that displays multiple reports that are relevant to the activities that you want to monitor. For more information about activities that you can perform on dashboards, see [About Dashboards](#).

 **Warning**

Infoblox recommends that you do not modify the predefined dashboards even if you have appropriate permissions. Editing the default dashboards changes the default settings and your changes become permanent. In addition, you might not be able to see the latest changes made by Infoblox. You can select a default dashboard and clone it to modify any of the settings, such as permissions, panels, and so on. For information about cloning dashboards, see [About Dashboards](#).

You can apply filters and view the dashboards in table, stacked area, or in both views. The following table lists the dashboard categories and their corresponding dashboards. The dashboards have been explained in detail in the later sections.

### Dashboard Categories

Dashboard Category	Corresponding Dashboard	Displays IDNs in Punycode (Yes/No)
<b>Audit Log Events</b>	See the following dashboard categories at <a href="#">Audit Log Events Dashboards</a>	
	For Audit Log Events	Yes
	For Audit Log WAPI Events	Yes
<b>IPAMv4 Utilization</b>	See the following dashboard categories at <a href="#">IPAMv4 Utilization Dashboards</a> :	
	DHCPv4 Top Utilized Networks	Yes
	DNS Statistics per DNS View	Yes
	DNS Statistics per Zone	Yes

Dashboard Category	Corresponding Dashboard	Displays IDNs in Punycode (Yes/No)
	IPAMv4 Network Usage Statistics	Yes
	IPAMv4 Network Usage Trend	Yes
	IPAMv4 Top Utilized Networks	Yes
	Managed DDI Peak IP Usage Trend	NA
	DNS Object Count Trend for Flex Grid License	NA
	IP Address Usage Report	NA
	VLAN Conflict	NA
	License Pool Utilization	No
	IPAM Prediction Dashboard	No
<b>Devices (Discovery)</b>	See the following dashboard categories at <a href="#">Devices (Discovery) Dashboards</a> :	
	Inactive IP Addresses	Yes
	Port Capacity Delta by Device	Yes
	Port Capacity Trend	Yes
	Port Capacity Utilization by Device	Yes
	IP Address Inventory	Yes
	Network Inventory	Yes
Network Insight	End Host History	Yes
	Device Interface Inventory	Yes
	Device Inventory	Yes
	Device Components	Yes
	IPAMv4 Device Networks	Yes
	Device Advisor	Yes

Dashboard Category	Corresponding Dashboard	Displays IDNs in Punycode (Yes/No)
<b>DHCP Dashboards</b>	See the following dashboard categories at <a href="#">DHCP Dashboards</a> :	
DHCP Fingerprints	Device Trend	Yes
	Device Class Trend	Yes
	Top Devices Identified	Yes
	Top Devices Denied an IP Address	Yes
	Top Device Classes	Yes
	Device Fingerprint Change Detected	Yes
DHCP Lease	DHCP Lease History	Yes
	DHCP Top Lease Clients	IDN is not supported
	<i>DHCP LPS Usage Report</i>	NA
DHCP Performance	DHCPv4 Range Utilization Trend	Yes
	<i>DHCPv4 Usage Trend</i>	Yes
	<i>DHCP Message Rate Trend</i>	Yes
	<i>DHCPv4 Usage Statistics</i>	No
<b>Cloud Dashboards</b>		
	<a href="#">VM Address History</a>	Yes
<b>System Utilization</b>	See the following dashboard categories at <a href="#">System Utilization Dashboards</a>	
	CPU Utilization Trend	Yes
	Memory Utilization Trend	Yes
	Traffic Rate by Member	Yes
	SPLA Grid Licensing Features Enabled	NA
	Managed DDI Features Enabled	Yes

Dashboard Category	Corresponding Dashboard	Displays IDNs in Punycode (Yes/No)
<b>DNS Dashboards</b>	See the following dashboard categories at <a href="#">DNS Dashboards</a>	
DDNS Query	DDNS Update Rate Trend	Yes
DNS Performance	DNS Response Latency Trend	Yes
	Managed DNS Peak Usage Trend	NA
	DNS QPS Usage Report	NA
DNS Query	DNS Top Requested Domain Names	Yes
	DNS Top Clients	Yes
	DNS Top Clients Per Domain	Yes
	DNS Query Rate by Query Type	Yes
	DNS Query Rate by Member	Yes
	DNS Replies Trend	Yes
	DNS Response Latency Trend	Yes
	DNS Top Clients Per Domain	Yes
	The DNS Top NXDOMAIN / NOERROR (no data)	Yes
	DNS Top SERVFAIL Errors Sent	Yes
	DNS Top SERVFAIL Errors Received	Yes
	DNS Top Timed-out Recursive Queries	Yes
	DNS Query Trend per IP Block Group	Yes
	DNS Effective Peak Usage Trend for SPLA Grid License	Yes
	DNS Daily Peak Hour Query Rate by Member	No
	DNS Daily Query Rate by Member	No
	DNS Cache Hit Rate Trend	No

Dashboard Category	Corresponding Dashboard	Displays IDNs in Punycode (Yes/No)
System Capacity	System Capacity Prediction Trend	No
<b>Internal Dashboards</b>	See the following dashboard categories at <a href="#">Internal Dashboards</a>	
	Reporting Index Usage Statistics	NA
	Reporting License Usage	No
	Reporting Volume Usage Trend per Category	NA
	Reporting Volume Usage Trend per Member	NA
License Dashboards	<a href="#">License Pool Utilization</a>	No
DNS Capture	DNS Domain Query Trend	Yes
	DNS Domains Queried by Client	Yes
	DNS Top Clients by Query Type	No
	DNS Top Clients Querying MX Records	No
DNS Traffic Control	DNS Traffic Control Resource Availability Status	No
	DNS Traffic Control Resource Availability Trend	No
	DNS Traffic Control Resource Pool Availability Status	No
	DNS Traffic Control Resource Pool Availability Trend	No
	DNS Traffic Control Resource SNMP Trend	No
	DNS Traffic Control Response Distribution Trend	No
DNS Reclamation	DNS Scavenged Object Count Trend	Yes
<b>Security Dashboards</b>	See the following dashboard categories at <a href="#">Security Dashboards</a>	
	FireEye Alerts	Yes
	DNS Top RPZ Hits	Yes
	DNS Top RPZ Hits by Clients	Yes

Dashboard Category	Corresponding Dashboard	Displays IDNs in Punycode (Yes/No)
	Threat Protection Event Count Dashboard	Yes
	Threat Protection Top Rules Logged	Yes
	DNS Top Tunneling Activity	Yes
	DNS Tunneling Traffic by Category	Yes
	Top Malware and DNS Tunneling Events by Client	Yes
	Detailed RPZ Violations by Subscriber ID	NA
	DNS RPZ Hits Trend by Mitigation Action	Yes
	DNS Firewall Executive Threat Report	No
	Top DNS Firewall Hits	No
	Malicious Activity by Client	Yes
<b>Ecosystem Dashboards</b>	See the following dashboard categories at <a href="#">Ecosystem Dashboards</a>	
	User Login History Report	Yes
	Subscription Data	Yes
	Publish Data	Yes



#### Note

To view RPZ-related reports or dashboards, you must enable RPZ logging on the **Logging** tab in the *Grid DNS Properties* editor.

## Home Dashboards

You can view the overall summary of DNS, DHCP, and IPAM activities in the **Home Dashboard** page. This page presents a summary view of the following:

- **DDI Summary:** Presents statistical information about the DNS, DHCP, and IPAM activities of all Grid members.
- **DNS:** Displays the statistical summary of DNS activities. You can export the search results, open in search, and refresh.
- **DHCP:** Displays the statistical summary of DHCP activities. You can export the search results, open in search, and refresh.
- **IPAM:** Displays the summary of the *Top 10 IPAMv4 Utilized Networks* dashboard.
- **Reporting Health:** You can view the license usage by the reporting server:
  - **Today's License Usage:** Current license usage by the indexer.

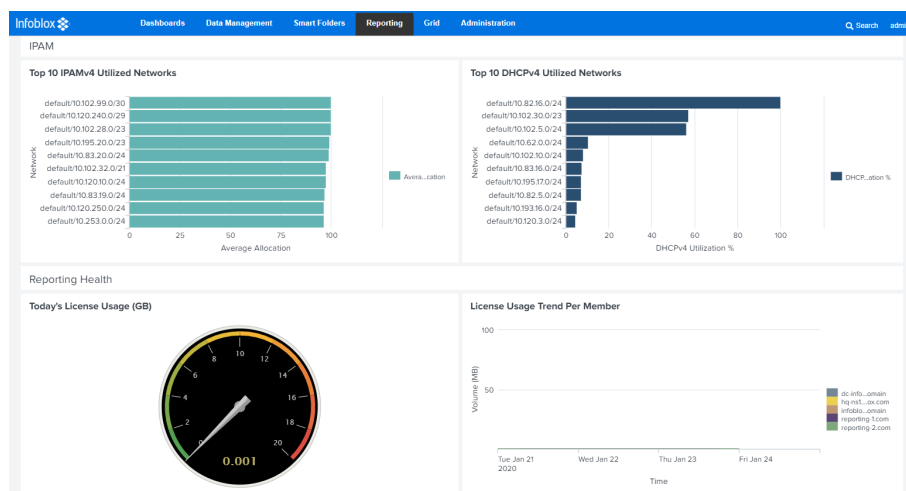
- **License Usage Trend per Member:** License usage by the indexer per member.



## Note

When you click **Open in Search** for a report, dashboard, or alert, the content of the entire page is encoded and displayed in the **Search** page. To avoid encoding, go to **Activity** tab -> **Jobs**. The **Jobs** page lists the search job history in the form of links. The top one is the latest search job executed by the alert or dashboard or report. The search string is not encoded when you click this link to run the search.

## Reporting Home Dashboard



On the Home Dashboard, you can also work with searches as described in the following sections:

- About Searches
- Best Practices for Customizing Searches
- Creating Reports from a Search
- Saving a Search as a Dashboard Panel
- Exporting Search Results
- Saving Search as Alerts

## About Searches

Searches are criteria that the reporting server uses to save reports and dashboard panels. Each predefined report has an associated search. For more information, refer to the official Splunk documentation: <http://docs.splunk.com/Documentation>.

To run a search:

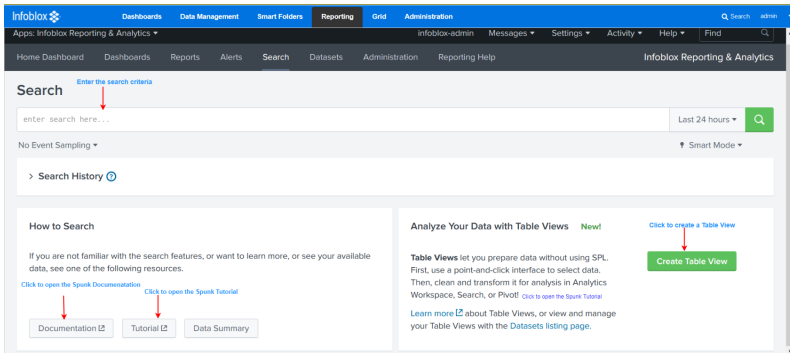
1. From the **Reporting** tab, select the **Search** tab.
2. Enter the search criteria. Use the auto-open search tips from Splunk.
3. If necessary, select a time range in the time range picker at the end of the search bar. By default, it is set to **Last 24 Hours**.
4. Click the Search icon.

The search results are based on the most seen events for the dashboards listed in the table below. To know more about dedup searches, reports, or dashboards, refer to <https://docs.splunk.com/Documentation/Splunk/8.2.4/SearchReference/Dedup>.

*Dashboards and Deduplicate Key(s)*

Dashboard	Deduplicate Key(s)
Inactive IP Addresses, for more information see <a href="#">Devices (Discovery) Dashboards</a> .	Network view + IP address
For more information about these dashboards, see <a href="#">IPAMv4 Utilization Dashboards</a> .	
DHCPv4 Top UtilizedNetworks	Network view + network
DNSStatistics per DNSView	DNS view
DNSStatisticsperZone	DNS view + DNS zone
IPAMv4 Network Usage Statistics	Network view + Network
IPAMv4 Top UtilizedNetworks	Network view + Network

### Sample Search Summary

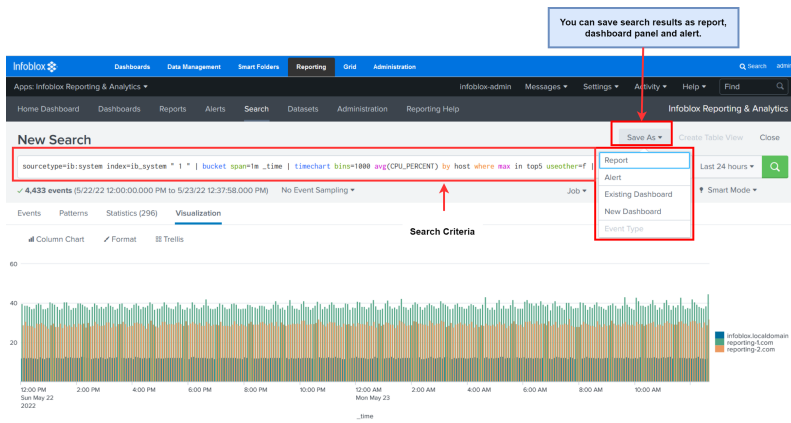


The search results are displayed in the **New Search** panel, as illustrated in the *New Search View*. In the **New Search** panel, you can save search results as **Report**, **Dashboard Panel**, and **Alert**.

When you deploy reporting clustering, we enable Splunk configuration to prevent data loss from forwarders, which may cause duplicated events in the indexer under certain circumstances. When you view reports and dashboards, the events that are already deduped are not duplicated again. However, if you view raw search events (such as write your own search against the indexed data directly), you may still see the duplicated events.

### New Search View





## Best Practices for Customizing Searches

You can optimize the performance of your reporting server and more efficiently view and manage your reports. Depending on the type of search and the data you want to search for, Infoblox recommends that you use the following guidelines:

- Specify shorter start and end times whenever possible.
  - Time range is one of the most important factors for search performance. Depending on the number of events that need to be loaded from the disk, it might take a long time when you specify a wider time range as it involves a large amount of data.
- Be specific about the fields you use.
  - Rare searches are faster than dense searches, so be specific whenever possible.
  - Start a search from a smaller dataset and then gradually apply it to bigger dataset.
  - When experimenting searches, start with a small date and time range, and then apply it to a bigger time range only when it is optimized.
- If a search is running for a long time, you can use the **Pause** and **Stop** buttons.
  - You can tune the search criteria and run it again if you stop an ongoing search job.
- Configure the panels to display data only if you have specific input instead of adding too many panels to the dashboard.
- Scheduling expensive searches.
  - You can configure reports and dashboards by scheduling searches because prefetched search results are displayed each time the reports and dashboards are opened. This reduces the workload on the reporting server without data freshness.
- Stagger scheduled searches.

Try to stagger your searches whenever possible. When you define how often the reporting server runs a search, be aware of other searches that the server is running. When you schedule the server to run many searches at the same time, the server performance can be negatively affected.

## Creating Reports from a Search

You can create reports by saving a search as a report. To save a search as a report:

1. From the **Reporting** tab, select the **Search** tab.
2. Enter the search criteria and then click the Search icon. The search results are displayed in the **New Search** panel.
3. From the **Save As** drop-down list, click **Report** to generate a report.
4. Enter title and description.
5. Click **Save**.
6. Do one of the following in the *Your report has been created* dialog box:
  - Click **View** to view your report on the Report page.
  - Click **Continue Editing** to edit.

- Click **Add to Dashboard** to add new report to the dashboard panel.

You can also complete the following settings in the *Your report has been created* dialog box:

- **Permissions:** Click this to edit permissions for your report, as described in Editing Permissions, see [Administrative Permissions](#).
- **Schedule:** Click this to schedule a report. For information about scheduling reports, see [About Reports](#).
- **Acceleration:** For more information, refer to the Splunk documentation.

## Saving a Search as a Dashboard Panel

You can save a search as a dashboard panel.

Do the following to save a search as a dashboard panel:

1. On the **Reporting** tab, select the **Search** tab.
2. Enter the search criteria and then click the Search icon. The search results are displayed in the **New Search** panel.
3. From the **Save As** drop-down list, choose **New Dashboard** to create a dashboard panel, or, you can choose **Existing Dashboard** to save the search to an existing dashboard panel.
4. In the *Save Panel to New Dashboard* dialog box, complete the following:
  - a. In the **Dashboard Title** field, enter a title.
  - b. Click **Edit ID** to modify the **Dashboard ID** field. It should only contain letters, numbers, and underscores.
  - c. In the **Description** field, type a description.
  - d. Select **Classic Dashboards**, or select **Dashboard Studio** and choose **Absolute** or **Grid** layout to create a dashboard.  
**Classic Dashboards** type of dashboard is the traditional splunk dashboard builder.  
**Dashboard Studio** type of dashboards are new type of dashboards available with the latest splunk version.
  - e. Click **Save to Dashboard**.
5. When prompted, click **View Dashboard** to view the dashboard in the **Dashboard** panel. For more information, see [About Dashboards](#).



### Note

There are no pre-defined dashboards available for the **Dashboard Studio**.

## Exporting Search Results

You can export the data in the selected search in CSV (comma separated value) or XML format. Note that this may take a long time depending on the amount of data you want to export. To schedule the export of search results to an FTP, SCP, or TFTP server configured on the **Set up** page, select **File Transfer Action** when creating a scheduled alert, as described in [Creating Scheduled Alerts](#), see [About Alerts](#).

To export data in a selected search:

1. From the **Reporting** tab, select the **Search** tab.
2. Enter the search criteria and then click the Search icon. The search results are displayed in the **New Search** panel.
3. Click the Export icon  
to export search results.
4. In the *Export Results* dialog box, complete the following:
  - **Format:** Select **CSV**, **XML** or **JSON** from the **Format** drop-down list.
  - **File Name:** Specify a file name for the export file. This is optional.
  - **Number of Results:(Limited or Unlimited).** If you select **Limited**, enter the number of results to be exported in the **Max Results** field.
5. Click **Export**.

## Saving Search as Alerts

To save a search as an alert:

1. From the **Reporting** tab, select the **Search** tab.
2. Enter the search criteria and then click the Search icon.
3. From the **Save As** drop-down list, click **Alert**.
4. In the **Save As Alert** dialog box, specify all alert settings. For information about creating scheduling alerts, see [About Alerts](#).
5. Click **Save**.

## About Alerts

You can configure alerts to trigger actions when certain events occur. When you set up an alert, search results trigger an alert action if they match the alert conditions. You can configure an alert to send an email notification, SNMP trap, and log a message in the syslog. Note that alerts are executed based on update frequencies for each corresponding search. For example, *DHCP Lease History* alerts are executed every 10 minutes, and Device Trend alerts are executed every 30 minutes at the 17th and 47th minutes of each hour (one minute after the search updates). For information about search indexes and update time intervals, see [About Reports](#). You can also throttle an alert if you want to change its frequency. For more information, refer to the Splunk documentation.

You can do the following in the **Alerts** page:

- Create scheduled alerts, as described in [Creating Scheduled Alerts](#) below.
- Edit permissions, as described in [Editing Permissions](#), see [Administrative Permissions](#).
- Edit alert type, trigger condition, and alert actions, as described in [Editing Alerts](#) below.
- Clone an alert, as described in [Cloning Alerts](#) below.

## Creating Scheduled Alerts

You can schedule an alert to notify when a scheduled report returns results that meet a specific condition. The appliance sends an alert when it encounters the trigger condition.

1. From the **Reporting** tab, select the **Alerts** tab -> select an *alert* and click **Open in Search**.
2. From the **Save As** drop-down list, click **Alert**.
3. In the *Save As Alert* dialog box, complete the following:
  - Specify the title and description.
  - **Alert Type**: Select **Scheduled**
  - **Time Range**: Specify the time range. For example, you can select **Run Every Day**.
  - **Schedule At**: Specify the time.
  - **Trigger Condition**: Specify trigger conditions. For more information, refer to the Splunk documentation.
  - **Trigger Actions**: Click this to configure alert actions. You can select the following:
    - **Send SNMP Trap**: Select this to enable SNMP traps. For information about how to trigger SNMP traps for reporting event types, see [Configuring SNMP](#).
    - **Send email**: Select this to send alert notification through email. You can specify email address in the **To** text box.
    - **Send to Syslog**. Select this to log a message in the syslog. If you configure this option with an alert, the message goes to the syslog on the reporting member or indexer.
    - **File Transfer Action**: Select this to upload the search results to an FTP or SCP or TFTP server configured on the **Set up** page. For information about Reporting (Index) Storage space, see [Grid Reporting Properties](#).
4. Click **Save**.

## Editing Alerts

You can edit alert type, trigger condition, and alert actions, as follows:

1. From the **Reporting** tab, select the **Alerts** tab -> select an *alert*.
2. From the **Edit** drop-down list, choose **Edit Alert** to edit the alert settings. In the *Edit Alert* dialog box, make the required changes.
3. Click **Save**.

## Cloning Alerts

1. From the **Reporting** tab, select the **Alerts** tab.
2. Select an alert you want to clone, click **Edit** -> **Clone**.
3. Enter a title and a description. Click **Clone Alert**.



### Note

Infoblox recommends that you save the cloned alerts only in the Infoblox Reporting & Analytics (infoblox) app.

## Configuring Email Notification Settings

You can enable the appliance to send email messages to specified recipients when the alert is triggered. You can configure email settings for alerts, scheduled reports, and scheduled PDF delivery.

To configure email properties for alerts and PDF delivery:

1. From the **Reporting** tab -> **Settings** tab -> click **Server settings**.
2. Click **Email settings**.
3. Specify the mail host. The default is local host.
4. Optionally, you can specify username and password.
5. Specify **Email Format**.
6. In the **Specify PDF Report** Settings, specify the paper size, paper orientation, and also the path to logo image.
7. Click **Save**.



### Note

You can configure email addresses when scheduling dashboard PDFs, scheduling reports, and creating alerts.

## About Reports

Infoblox provides reports that are named by core network service functions, such as DNS query and system utilization. Reports contain predefined search criteria that retrieve specific data from the reporting database. Each report is associated with a search. It is not recommended to modify predefined reports. However, when you run a search, you can save it as reports and share it with other users. You can also create a new report by cloning an existing report, and then modify the search criteria.

You can also create a personal report in two different ways:

- Clone a report, as described in the [Cloning Reports](#) section below.
- Saving a search as a report, as described in the [Creating Reports from a Search](#) section below.



### Note

IDNs are not supported on the reporting server. The reporting server manages IDNs in punycode. The reports generated by collecting reporting data from the DNS server displays all the data in punycode only.

When you upgrade to 7.3.x, all NIOS Global reports are migrated to the **Dashboards** panel without filters. However, the filter conditions configured in NIOS are reflected in the **Dashboards** panel. All NIOS System and Global searches are migrated to the **Reports** panel.

You can do the following in the **Reports** tab:

1. From the **Reporting** tab, select the **Reports** tab -> select a report.
2. In the **Reports** panel, you can do the following:
  - Open a report in the **Search** page and edit the report using the **Save As** menu. For information about Searches, see [Home Dashboards](#).
  - You can do the following using the **Edit** drop-down list:
    - Edit Description
    - Edit Permissions, as described in the Editing Permissions section, see [Administrative Permissions](#).
    - Edit Schedule, as described in the Scheduling Reports section below.
    - Clone a report, as described in the Cloning Reports section below.

## Reporting Indexes and Update Time Intervals

The Reporting Indexes table below lists the search indexes that the reporting server uses to generate reports. It contains information about the frequency of the summary report updates for each report and the percentage of the total index space allocated for each report category. Use this information to plan your reporting strategy for the Grid so you can optimize the performance of the reporting server.

Each summary report or search has its own update frequency. For example, the *DNS Top Requested Domain* report updates its data every 30 minutes, starting at the 4<sup>th</sup> minute of each half hour. It collects report data during the first 30 minutes of the previous 60 minutes. For example, if the report starts an update at 6:04 a.m., the data it collects is from 5:04 a.m. to 5:34 a.m.

The reporting server also uses this information to generate alerts. For example, once configured, Top Devices Identified alerts are executed at the 17<sup>th</sup> and 47<sup>th</sup> minutes of each hour (one minute after each update), regardless of whether DHCP fingerprint detection is enabled or disabled. For information about alerts, see [About Alerts](#)



### Note

The maximum retention period for the reporting data is 136 years. However, the data is removed from the database if the data exceeds the maximum limit for a reporting index and when the data crosses the retention period (after 136 years).

### Reporting Indexes

Indexes	Reports/Searches	Summary Report Data Updates	Default Maximum Index Size (% of Total Index Storage)
Device (Discovery)			0%
	Inactive IP Addresses	N/A	
	Port Capacity Utilization by Device	N/A	

Indexes	Reports/Searches	Summary Report Data Updates	Default Maximum Index Size (% of Total Index Storage)	
	Port Capacity Trend	N/A		
	Port Capacity Delta by Device	N/A		
	IP Address Inventory	N/A		
	Network Inventory	N/A		
	IPAMv4 Device Networks	N/A		
	Device Interface Inventory	N/A		
	Device Inventory	N/A		
	Device Components	N/A		
	Device Advisor	N/A		
	End Host History	N/A	0%	
IPAMv4			5%	
	IPAMv4 Network Usage Statistics (Detailed)	N/A		
	DNS Statistics per DNS View (Detailed)	N/A		
	DNS Statistics per Zone (Detailed)	N/A		
	IPAMv4 Top Utilized Networks (Detailed)	N/A		
	DNS Object Count Trend for Flex Grid License	Data is generated once every 24 hours		

Indexes	Reports/Searches	Summary Report Data Updates	Default Maximum Index Size (% of Total Index Storage)	
DNS			10%	
	DDNS Update Rate Trend	N/A		
	DNS Response Latency Trend (Summary)	N/A		
	DNS Cache Hit Rate Trend	N/A		
	DNS Query Rate by Query Type	N/A		
	DNS Query Rate by Server (Detailed)	N/A		
	DNS Replies Trend	N/A		
	DNS Query Trend Per IP Block Group	N/A		
	FireEye Alerts	N/A		
DNS Summary			10%	
	DDNS Update Rate Trend (Summary)	Every 30 minutes, starting at the 6 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DNS Response Latency Trend (Summary)	Every 30 minutes, starting at the 20 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DNS Top Requested Domain Names (Summary)	Every 30 minutes, starting at the 4 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DNS Cache Hit Rate Trend (Summary)	Every 30 minutes, starting at the 8 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DNS Effective Peak Usage Trend for SPLA Grid License	NA		

Indexes	Reports/Searches	Summary Report Data Updates	Default Maximum Index Size (% of Total Index Storage)	
	DNS Top Clients (Summary)	Every 30 minutes, starting at the 2 <sup>nd</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DNS Query Rate by Query Type (Summary)	Every 30 minutes, starting at the 10 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DNS Query Rate by Member (Summary)	Every 30 minutes, starting at the 12 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DNS Daily Query Rate by Member (Summary)	Every day, starting at 00:32 each day. Data covers from 00:00 of yesterday to 00:00 of today.		
	DNS Daily Peak Hour Query Rate by Member (Summary)	Every 60 minutes, starting at the 34 <sup>th</sup> minute of each hour. Data covers from the top of last hour to the top of current hour.		
	DNS Replies Trend (Summary)	Every 30 minutes, starting at the 18 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DNS Top Clients Per Domain (Summary)	Every 30 minutes, starting at the 3 <sup>rd</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DNS Top NXDOMAIN / NOERROR (no data) (Summary)	Every 30 minutes, starting at the 5 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DNS Top SERVFAIL Errors Sent (Summary)	Every 30 minutes, starting at the 6 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DNS Top SERVFAIL Errors Received (Summary)	Every 30 minutes, starting at the 7 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DNS Top Timed-Out Recursive Queries (Summary)	Every 30 minutes, starting at the 8 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		



Indexes	Reports/Searches	Summary Report Data Updates	Default Maximum Index Size (% of Total Index Storage)	
	DNS Top RPZ Hits (Summary)	Every 10 minutes, starting at the 2 <sup>nd</sup> minute of each 10 minute. Data covers from 20 minute ago to 10 minute ago.		
	DNS Top RPZ Hits by Clients (Summary)	Every 10 minutes, starting at the 2 <sup>nd</sup> minute of each 10 minute. Data covers from 20 minute ago to 10 minute ago.		
	DNS Scavenged Object Count Trend	Every 30 minutes, starting at every 21 <sup>st</sup> and 51 <sup>st</sup> minute of each hour.		
	DHCP			2 Months
		DHCPv4 Usage Statistics (Detailed)	N/A	
		DHCPv4 Range Utilization Trend (Summary)	N/A	
		DHCP Message Rate Trend (Detailed)	N/A	

Indexes	Reports/Searches	Summary Report Data Updates	Default Maximum Index Size (% of Total Index Storage)	
DHCP Summary			24 Months	
	Device Trend (Summary)	Every 30 minutes, starting at every 16 <sup>th</sup> and 46 <sup>th</sup> minutes of each hour. Data covers the first 30 minutes of the previous 60 minutes.		
	Device Class Trend (Summary)	Every 30 minutes, starting at every 16 <sup>th</sup> and 46 <sup>th</sup> minutes of each hour. Data covers the first 30 minutes of the previous 60 minutes.		
	Top Devices Identified (Summary)	Every 30 minutes, starting at every 16 <sup>th</sup> and 46 <sup>th</sup> minutes of each hour. Data covers the first 30 minutes of the previous 60 minutes.		
	Top Devices Denied an IP Address (Summary)	Every 30 minutes, starting at every 19 <sup>th</sup> and 49 <sup>th</sup> minutes of each hour. Data covers the first 30 minutes of the previous 60 minutes.		
	Top Device Classes (Summary)	Every 30 minutes, starting at every 16 <sup>th</sup> and 46 <sup>th</sup> minutes of each hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DHCP Top Lease Clients (Summary)	Every 30 minutes, starting at the 16 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	DHCPv4 Range Utilization Trend (Summary)	Every 8 hours, starting at the 24 <sup>th</sup> minute of each half hour. Data covers the first 8 hours of the previous 8.25 hours.		
	DHCPv4 Usage Trend (Summary)	Every 8 hours, starting at the 22 <sup>nd</sup> minute of each half hour. Data covers the first 8 hours of the previous 8.25 hours.		
	DHCP Message Rate Trend (Summary)	Every 30 minutes, starting at the 14 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
DHCP Lease History			39%	
	DHCP Lease History (Detailed)	N/A		

Indexes	Reports/Searches	Summary Report Data Updates	Default Maximum Index Size (% of Total Index Storage)	
	Device Fingerprint Change Detected (Detailed)	N/A		
Security			1%	
	Threat Protection Event Count By Severity Trend (Summary)	N/A		
	Threat Protection Event Count By Member Trend (Summary)	N/A		
	Threat Protection Event Count By Rule (Summary)	N/A		
	Threat Protection Event Count By Category (Summary)	N/A		
	Threat Protection Event Count By Time (Summary)	N/A		
	Threat Protection Event Count By Member (Summary)	N/A		
	Threat Protection Top Rules Logged (Summary)	N/A		
	Threat Protection Top Rules Logged by IP (Summary)	N/A		
	DNS Top Tunneling Activity (Summary)	Every 30 minutes, starting at every 11 <sup>th</sup> and 41 <sup>st</sup> minute of each hour.		
	DNS Tunneling Traffic by Category (Summary)	Every 30 minutes, starting at every 11 <sup>th</sup> and 41 <sup>st</sup> minute of each hour.		
	Top Malware and DNS Tunneling Events by Client (Summary)	Every 30 minutes, starting at every 11 <sup>th</sup> and 41 <sup>st</sup> minute of each hour.		
	Cloud	VM Address History (Detailed)	N/A	0%
Audit Log	Audit Log Events (Detailed)	N/A	0%	

Indexes	Reports/Searches	Summary Report Data Updates	Default Maximum Index Size (% of Total Index Storage)	
	Audit Log WAPI Events	N/A	0%	
Ecosystem	Ecosystem Subscription Ecosystem Publication	N/A		
License	License Pool Utilization	N/A	1%	
System Utilization			15%	
	Memory Utilization Trend (Summary)	Every 30 minutes, starting at the 26 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes		
	Memory Utilization Trend (Detailed)	N/A		
	Traffic Rate (Detailed)	N/A		
	Traffic Rate by Member (Summary)	Every 30 minutes, starting at the 28 <sup>th</sup> minute of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	CPU Utilization Trend (Summary)	Every 30 minutes, starting at the top of each half hour. Data covers the first 30 minutes of the previous 60 minutes.		
	SPLA Grid Licensing Features Enabled	Every 24 hours for all IB-FLEX members in the Grid.		



When you filter a dashboard by a time frame that is larger than the maximum retention period, the reporting server returns data within the maximum retention period. For example, when you try to view data of the *CPU Utilization Trend* report for the past six months, the server only returns data up to the last two months.

## Cloning Reports

1. From the **Reporting** tab, select the **Reports** tab.
2. Select the report you want to modify, click **Edit -> Clone**.
3. Enter a new title and description.
4. Set its permissions. Select **Private** if you do not want to share the cloned report with other users. Select **Clone** if you want the cloned report to have the same permissions as the original report.
5. Click **Clone Report**.
6. Optionally, you can do the following:
  - Edit the permissions as described in the [Editing Permissions](#) section, see [Administrative Permissions](#).
  - Click **View** to view the cloned report.
  - Click **Add to Dashboard** to add the cloned report to the dashboard.

Click **Open in Search** to open the cloned report in the **Search** page.



#### Note

Infoblox recommends that you save the cloned reports only in the Infoblox Reporting & Analytics (infoblox) app.

## Deleting Reporting Data

To selectively delete reporting data from NIOS, you must first enable the `delete` permission for the local admin with superuser permission by using the `set reporting_user_capabilities` command in the NIOS CLI. Then, complete the following:

1. From the **Reporting** tab, select the **Search** tab.
2. Enter the search criteria in the search bar that returns the events you want to delete and click the **Search** icon.
3. After you confirm that the search results contain only those events that you want to delete, pipe the search to the `delete` command.

Example:

```
<Splunk_Query> | delete
```

For more information, see the Splunk documentation.



#### Notes

- You cannot retrieve the data once it is deleted.
- You cannot visualize the deleted data.
- The deleted data does not reduce any disk space.
- Frequent deletion of data may affect the search performance.
- You must not delete any of the `ib_threatdb*` index files as it results in loss of threat events data.

## Scheduling Reports

You can schedule a report to run on a scheduled interval and trigger an action each time it runs. When scheduling a report, you can set up an action to send an email to receive report results. In addition, you can export results in CSV (comma separated value) or XML format.

To schedule a report:

1. From the **Reporting** tab, select the **Reports** tab.
2. Select the report you want to schedule, click **Edit** -> **Edit Schedule**.  
Note that you can schedule a report when you save search results as reports. When you set the paper size to **A5**, the logo image and report name may overlap in the footer of the downloaded reports or reports sent through email.
3. In the *Edit Schedule* dialog box, select the **Schedule Report** checkbox.
4. Enter the **Schedule** and **Time Range**. For more information about how to use the **Schedule** and **Time Range** options, refer to the Splunk documentation.
5. Click **Add Actions**, and then to set an action for the scheduled reports, select **Send Email** to send an email that contains report results in text format, or as CSV or PDF attachments to a specified set of recipients.
  - Enter the email address in the **To** text box. To send the email message to multiple recipients, type a comma between email addresses.*Edit Schedule*

The screenshot shows the 'Edit Schedule' dialog box. At the top, there is a '+ Add Actions' button. Below it, a list of actions is shown, with 'Send email' selected and highlighted by a red box. The 'Send email' configuration includes:

- To:** test@gmail.com (highlighted by a red box)
- Priority:** Normal
- Subject:** Splunk Alert: \$name\$
- Message:** Default
- Include:**
  - Link to Report
  - Link to Results
  - Search String
  - Inline Table
  - Attach CSV
  - Attach PDF
  - Allow Empty Attachment
- Type:** HTML & Plain Text

At the bottom right, there are 'Cancel' and 'Save' buttons.

- In the **Include** section, select one of the following: **Inline Table**, **Attach CSV**, **Attach PDF**. Selecting **Attach PDF** or **Attach CSV** attaches the results of the report in the form of a CSV file or a PDF. You can also select the **Allow Empty Attachment** for attachment-based include options. Ensure that you specify this information.

6. Click **Save**.

Infoblox recommends that you do not select Link to Report, Link to Results, Search String in the Include section. These links might not work in your environment. Do not select the Run a Script option because there is no script to run.

## Configuring Logo Image in PDF Reports

All reports display the Infoblox logo by default. You can customize reports by removing the Infoblox logo, or by replacing it with your own company logo. The reporting server uses the latest image file that you have uploaded. Make sure that you upload a logo file that is in PNG format and has a file size that is smaller than 500 KB. Note that the image file name must be *pdf\_logo\_image.png*. Do not change the logo image file name. You can configure your PDF reports and schedule to send them through emails. For information about scheduling the delivery of reports, see [Scheduling Reports](#) below.



#### Note

In the footer of the report, you can view the logo image (if uploaded), panel name, and the timestamp when the report was downloaded. When there is no data in a single panel report, the downloaded PDF displays "No Results Found" along with "Last Updated" information. However, a report with multiple panels, displays only the panel name for the panel that does not have any data.

To upload logo image:

1. From the **Administration** tab, select the **Reporting** tab -> expand the Toolbar and click **Grid Reporting Properties**.  
or  
From the **Grid** tab, select the **Grid Manager** tab and click the **Services** tab. On the **Services** tab, select the **Reporting** tab and click **Edit** -> **Grid Reporting Properties** from the Toolbar.
2. In the Grid Reporting Properties editor, select the **PDF** tab and complete the following:  
**Logo Image:** Click **Upload** to open the *Upload* dialog box. Click **Select** to navigate to where the image file is located and click **Open**. Click **Upload** to upload the file. The appliance displays a preview of how it will appear on reports after you successfully upload the logo file. You can click **Clear** to remove the logo and upload a new one. Make sure that the file format and size meet the requirements; otherwise, the appliance displays an error message. You can click **Clear** to remove the uploaded image file.
3. Click **Save & Close**.

## About Dashboards

Dashboards provide summary views for most of the data and trends in your Grid. Grid Manager includes a list of default dashboards. For more information, see [Predefined Dashboards](#). Infoblox recommends not to modify the default dashboards. If you need to modify the settings of a default dashboard, you can either clone a default dashboard, or create a new dashboard from scratch and then add panels and reports. For example, you can create a new dashboard called "DNS and DHCP Activities," and then add DNS report, add DHCP related reports, such as DHCP Top Lease Clients and DHCP Lease History, to the new dashboard. When you save the "DNS and DHCP Activities" dashboard, the reporting server saves all the reports added to the dashboard and displays dashboard with updated data. By doing this, user-defined dashboards can provide single point of access to review multiple reports that are relevant to the activities you want to monitor. If you modify a default dashboard, you can reset to its default settings. Each dashboard comes with a set of filters to further refine report data.

When you upgrade to 7.3.0 or later, all your system reports are migrated to **Reporting** tab -> **Dashboards**.



#### Warning

Infoblox recommends that you do not modify the predefined dashboards even if you have appropriate permissions. Editing the default dashboards changes the default settings and your changes become permanent. In addition, you might not be able to see the latest changes made by Infoblox. You can select a default dashboard and clone it to modify any of the settings such as permissions, panels, and so on. For more information, see [Cloning Dashboards](#) below.

In the **Dashboard** panel, you can do the following using the **Edit** drop-down list:

- Edit panels as described in the [Editing Dashboards](#) section below.
- Edit source, title, and description.
- Edit permissions as described in the [Editing Permissions](#) section below.
- Clone a dashboard as described in the [Cloning Dashboards](#) section below.
- Reset dashboards, as described in the [Resetting Dashboards](#) section below.



### Note

You cannot set any dashboard as the home dashboard. The **Set as Home Dashboard** option available in the **Edit** drop-down list does not add a dashboard to the **Home Dashboard** tab.

## Creating New Dashboards

When you add a new dashboard, Grid Manager displays it in the **Reporting** -> **Dashboards** tab. You can add multiple panels and reports to the new dashboard.

To create a new dashboard:

1. From the **Reporting** tab -> select the **Dashboards** tab.
2. Click **Create New Dashboard**.
3. Complete the following:
  - **Dashboard Title:** Enter the dashboard title.
  - **Description:** Enter the dashboard description.
  - **Permissions:** Click **Shared in App** to share a dashboard to other users. Depending on their permissions, other users can edit the dashboard. When **Private** is selected, the dashboard is available only to the user who creates it. You can change permissions later while editing a dashboard.
  - Select **Classic Dashboards** or **Dashboard Studio** and choose **Absolute or Grid layout to create a dashboard in Dashboard Studio**.  
Classic Dashboards type of dashboard is the traditional Splunk dashboard builder.  
Dashboard Studio type of dashboards are new type of dashboards available with latest version of Splunk.  
For details on Splunk Dashboard Studio, please refer <https://docs.splunk.com/Documentation/Splunk/8.2.4/DashStudio/IntroFrame>.
4. Click **Create**.



### Note

- Uploading a background image or icons to **Dashboard Studio** type dashboards is not supported.
- There are no pre-defined dashboards available for **Dashboard Studio**.

## Cloning Dashboards

It is not recommended to modify default dashboards. You can select a default dashboard and clone it to modify any of the settings, such as permissions, panels, and so on.



### Note

You do not need any permission to create, modify, and delete your own personal dashboard. However, limited-access users need **Read** and **Write** permissions to modify cloned dashboards. For information about administrative permissions, see [Administrative Permissions](#).

When you clone a dashboard, you can do the following:

- View and set permissions, as described in [Editing Permissions](#), see [Administrative Permissions](#).
- Schedule PDF delivery, as described in [Scheduling PDF Delivery for Dashboards](#) below.
- Edit panels, as described in [Editing Dashboards](#) below.

To create a personal dashboard:



1. From the **Reporting** tab, select the **Dashboards** tab.
2. Select the dashboard you want to modify, click **Edit** -> **Clone**.
3. Enter a new title, ID, and description.
4. Set its permissions. Select **Private** if you do not want to share the cloned dashboard with other users.  
Select **Clone** if you want the cloned dashboard to have the same permissions as the original dashboard.
5. Click **Clone Dashboard**.
6. Optionally, you can edit permissions, as described in Editing Permissions, see [Administrative Permissions](#) or click **View** to view the cloned dashboard.

## Resetting Dashboards

Infoblox recommends not to modify default dashboards. However, when you make changes to the default dashboards, you can reset to its default settings.

To reset a dashboard:

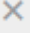
1. From the **Reporting** tab -> select the **Administration** tab.
2. Click **Reset Dashboards**.
3. Select the checkbox of the dashboard or Select all to select all the dashboards.
4. Click **Reset selected dashboards**.

The dashboard you have modified will reset to its default settings.

## Editing Dashboards

To edit the panels and filters of a dashboard, it is recommended to clone the default dashboard, and then add panels, filters and reports to the cloned dashboard. When you add a report to the panel, Grid Manager generates the corresponding dashboard in the panel. When you save the dashboard, Grid Manager updates reports in each panel. Alternatively, you can edit the XML source code to add filters and panels to a cloned dashboard, as described in Editing the XML Source Code of a Dashboard below.

To add panels and filters to a dashboard:

1. From the **Reporting** tab -> select the **Dashboards** tab -> select a *dashboard*.
2. From the **Edit** drop-down list, select **Edit Panels**.
3. In the **Edit: <Dashboard>** pane, you can click **Add Panel** or **Add Input** or **Edit Source**.  
Note that you cannot modify or delete the default values set for the dashboard filters. For example, you cannot delete or modify the filter All set for Members. When you add a new input using the editor, make sure that you edit the source and refer to the token for the input in the search string. By doing so, the search is updated when you change the input value. For information editing source, refer to Splunk documentation.
4. Optionally, you can click  to delete a filter. When you delete a filter, make sure that you delete the filter information from the XML source code as well. For information, see Editing the XML Source Code of a Dashboard below.
5. Expand the panel categories and select the panel you want to add. For detailed information about how to add panel categories, refer to the Splunk documentation.
6. Click **Add to Dashboard**.

## Editing the XML Source Code of a Dashboard



### Note

Before editing dashboards, forms, and panel files in simple XML source code, you should be familiar with the basic layout of dashboards and the XML elements that define them. Infoblox recommends that you save a copy of the source code of the dashboard before making any modifications.

To edit the XML source code for a dashboard:

1. From the **Reporting** tab, select the **Dashboards** tab.
2. Select a dashboard you want to edit the XML source code, click **Edit** -> **Edit Source**.

You can add filters, such as checkbox, drop-down list, radio button, and text box. For information about how to edit the XML source code, refer to the Splunk documentation.

#### Example - Adding an extensible attribute filter

If your reporting data contains the "location" extensible attributes associated to members, adding the following sample XML code to the XML source code will create an extensible attribute filter, **Member Location**:

```
<input type="dropdown" token="ea_location">
  <label>Member Location</label>
  <choice value="All">All</choice>
  <default>All</default>
  <search>
    <query>| inputlookup __grouping_by_ea_tag_lookup
      | spath input=EA path=Location output=EA_Location
      | stats count by EA_Location </query>
    <earliest>$time.earliest$</earliest>
    <latest>$time.latest$</latest>
  </search>
  <fieldForLabel>EA_Location</fieldForLabel>
  <fieldForValue>EA_Location</fieldForValue>
  <change>
    <condition value="All">
      <set token="ea_location_str"> | noop </set>
    </condition>
    <condition value="*">
      <set token="ea_location_str"> | spath input=EA path=Location
output=EA_Location
      | where EA_Location="$value$" </set>
    </condition>
  </change>
```

```
</input>
```

```
<search id="base_search">
```

```
<query>index=ib_system_summary report=si_cpu_usage
```

```
$members$
```

```
$ea_site_str$
```

```
$ea_location_str$
```

```
$group_by_str$
```

```
$group_by_stats$
```

```
    | timechart bins=1000 $calculation_method$(CPU_PERCENT) by  
$time_chart_field$
```

```
where max in $topn$ useother=f
```

```
    | interpolate 1200</query>
```

```
<earliest>$time.earliest$</earliest>
```

```
<latest>$time.latest$</latest>
```

```
</search>
```

## Scheduling PDF Delivery for Dashboards

To schedule PDF delivery for dashboards, you must first create a new dashboard. Ensure that email notification settings are configured prior to scheduling PDF delivery. For more information about Configuring Email Notification Settings, see [About Alerts](#). To schedule PDF delivery, you can use the dashboard **Edit** drop-down list. You can access the **Edit** drop-down list directly from a dashboard or from the **Dashboards** page.



### Note

Scheduled PDF delivery is not available for dashboards that include forms.

Do the following if the **Schedule PDF Delivery** option is disabled:

1. Open the dashboard for which you want to schedule PDF delivery.
2. Click **Open in Search** icon available at the bottom of the dashboard panel.
3. From the **Save As** menu, click **Dashboard Panel**.

To set up PDF delivery for the dashboard with a single panel:

1. From the **Reporting** tab, select the **Dashboards** tab.
2. Do one of the following:
  - Select the dashboard you want to schedule, click **Edit** -> **Schedule PDF Delivery**. If the **Schedule PDF Delivery** option is disabled, follow the steps as described above.

- Open the dashboard in the **Dashboards** page and click **Schedule PDF Delivery** from the **Edit** drop-down list.
3. In the *Edit PDF Schedule* dialog box, do the following:
- Select the **Schedule PDF Delivery** checkbox to enable PDF delivery.
  - Select a schedule. For more information, refer to the Splunk documentation.
  - In the **Email To** text box, specify email address.
  - Select paper size and paper layout. You can change the paper size and paper layout, if data is not displayed properly in the PDF delivery.
  - To receive dashboard PDFs immediately, click **Send Test Email**.

4. Click **Save**.



Note

To set up PDF delivery for the dashboard with multiple panels, repeat the above steps from step 1 to step 4 and add other panels to the dashboard created for the first panel.

*Edit PDF Schedule dialog box*

Click **Send Test Mail** to receive Dashboard PDFs Immediately.

## About Dashboard Filters

You can apply different filters to control the data displayed in the dashboards. The data on the dashboard is displayed based on the various filter criteria you select.

To apply a filter:

1. From the **Reporting** tab, select the **Dashboards** tab -> select a *dashboard*.
2. Apply filter criteria appropriately and click **Submit**.

The dashboards display results based on the filters that you apply.

The most common filters are as follows:

- **Time**
- **Top N**: Topmost filter options. The default is 10. You can select from a set of fixed values for the TopN filter setting: 5, 10, 20, 50, 100, 200, 250, or 500.
- **Members**: Grid members configured on the appliance.
- **Network**
- **Member Site**, as described in Applying Extensible Attribute Filters below.

## Applying Time Filters

You can generate a dashboard for a specific time interval by applying time filters. You can filter results by preset time ranges, create custom time ranges, specify time ranges based on dates or date and time, or work with advanced features in the time range picker. For information about Time range picker, refer to the Splunk documentation.

The date and time displayed in the **Time** filters are based on the time zone set in your user profile by default. For more information about how to configure a time zone, see Setting the Browser Time Zone in [Setting Login Options](#).

However, the timestamp displayed in the results for a dashboard is based on the time zone configured on the reporting server.



### Note

The NIOS reporting data is updated at a certain time interval, rather than updating continuously. Therefore, the **Real-time** option in the **Time** filter might not work for most of the dashboards. For information about update time intervals, see Reporting Indexes and Update Time Intervals in [About Reports](#).

## Applying Extensible Attribute Filters

You can use extensible attribute filters to narrow down the search by including only members that contain certain extensible attribute values. An extensible attribute added to a member is displayed in the **Extensible Attribute** filter. For information about managing extensible attributes, see [Managing Extensible Attributes](#). When you configure group-by-extensible-attribute search and apply the **Extensible Attribute** filter, the dashboard displays results for grouped members that have the same extensible attribute value for the **Site** extensible attribute. If you have configured multiple attribute values for a member, then applying the **Extensible Attribute** filter displays all the attribute values associated with that member. For example, if member 1 has predefined attribute **Site** with attribute values `member a` and `member b` and member 2 has predefined attribute **Site** with attributes values `member c` and `member d`, then the dashboard displays `member a` and `member b` when you apply the **Member Site <member 1>** filter.

In addition, you can apply **Group By EA Tag** filter and group members with the same extensible attribute value so that instead of displaying data per member, the reports display data per group of members with the same value for the **Site** extensible attribute. When you apply **Group By EA Tag** filter, you can set the data calculation method to decide which statistic value [Aggregate, Average, or Maximum] you want to be displayed for grouped members. You can group by Active Directory Sites for the *IPAMv4 Network Usage Statistics* report, *IPAMv4 Top Utilized Networks* report, and *DHCPv4 Network Usage Statistics* report.

To apply an extensible attribute and group by EA tag/field filters:

1. From the **Reporting** tab, select the **Dashboards** tab -> select a dashboard.
2. In the filter section, complete the following:
  - **Member <Extensible Attribute>**: Select an extensible attribute configured for a member. If you need an additional extensible attribute filter, you must first clone the default dashboard, and then add an extensible attribute filter by editing the XML source code. For information, see [Editing the XML Source Code of a Dashboard](#) above.
  - **Group By EA Tag/Field**: Select an extensible attribute to enable the reporting server to group networks by members that have certain extensible attribute tags or fields. Note that this option is available for specific dashboards only.  
 Note that if you use special characters in the extensible attribute name, the appliance replaces these special characters with equivalent values. For example, the extensible attribute **Site In London** is displayed as Site20In20London in the **Group By EA Tag/Field** drop-down list. In this example, space is replaced with 20. If you add the extensible attribute **London@**, it is displayed as **London40** in the **Group By EA Tag/Field** drop-down list.
  - **Calculation Methods**: This field is enabled only when you select the **Group by EA Tag/Field** checkbox. The displayed result varies based on your search definitions. The result values can contain information such as event counts, DNS queries, traffic rate, and usage trends. For example, when you select **Maximum**, the *DNS Query Rate by Member* dashboard shows all the members that have the same extensible attributes and members with the maximum DNS queries, and the *Threat Protection Event Count By Member* dashboard shows the members that have the same extensible attributes and maximum event counts. Select one of the following methods:
    - **Aggregate**: Displays the sum of values for individual members in a group.
    - **Average**: Displays the mathematical average of a group. This value is obtained by adding values for all members in a group and then dividing the total by the number of members.
    - **Maximum**: Displays the maximum value among the members in a group.



#### Note

When you apply **GroupByEATag/Field** in Active Directory Sites supported reports, the values displayed in these reports are aggregated sum of absolute values (sum of values of individual networks in a group) and utilization% is the mathematical average of the group.

You can configure the group-by-extensible-attribute filter and data calculation methods for the following dashboards only. For more information about these dashboards, see [Predefined Dashboards](#).

- CPU Utilization Trend (Detailed)
- IPAMv4 Network Usage Statistics
- IPAMv4 Top Utilized Networks
- IPAMv4 Network Usage Trend
- DDNS Update Rate Trend
- DHCPv4 Usage Statistics
- DNS Daily Query Rate by Member
- DNS Query Rate by Member
- DNS Daily Peak Hour Query Rate by Member
- DNS Response Latency Trend
- DNS Cache Hit Rate Trend
- DNS Traffic Control Resource Availability Trend
- DNS Traffic Control Resource Pool Availability Trend
- DNS Traffic Control Response Distribution Trend
- Memory Utilization Trend
- Traffic Rate by Member
- Threat Protection Event Count By Member
- Threat Protection Event Count By Member Trend

## Reporting Data Model

This section contains information about fields that are included in the reports and dashboards. You can find the commonly extracted fields and their specifications such as data source and range, which can help you better define your dashboards and searches.

### Splunk default fields

Splunk server adds the following default fields to each event in every index.

Field Name	Description	Values/Range
date_hour	Indicates the hour when an event occurred. To narrow your search for specific event timestamps, you can use the default datetime fields. Click <a href="#">here</a> for more information on datetime fields.	Range: 0-23
date_mday	Indicates the day of the month when the event occurred	Range: 1-31
date_minute	Indicates the exact minute when the event occurred	Range: 0-59
date_month	Indicates the month during which an event occurred	
date_second	Indicates the second in which an event occurred	Range: 0-59
date_wday	Indicates the day of the week in which an event occurred	Example: Sunday, Monday, etc.
date_year	Indicates the year in which an event occurred	
date_zone	Indicates the time for the local timezone of an event, expressed as hours in Unix Time	
eventtype	Indicates events of the same type based on a given search. Click <a href="#">here</a> for more information.	Example: splunkd-log
host	Contains information about the originating hostname or a network IP address that generates the event	Example: reporting-1.com
index	Contains the name of the index with which a given event is indexed	Example: ib_dns_summary
linecount	Contains information about the number of lines in an event before it is indexed	Example: 1

punct	Contains information about the pattern of the first thirty punctuation characters in the first line of the event with which it is associated. It shows how an event looks when all letters, numbers, and spaces are removed and contains characters such as periods, colons, parentheses, quotes, question marks, dashes, and underscores. Click <a href="#">here</a> for more information.	Example: -_:::\N:____/_=
source	Contains the name of the file, stream, or other input details from which the event originates	Example: si-search-dns-query-reply
sourcetype	Specifies the format of data input from which the event originates	Stash
splunk_server	Contains the name of the Splunk server that comprises the event	Example: reporting-2.com-2- <secondary server>
splunk_server_group	Contains the name of the Splunk server group	String

## Commonly Extracted Fields

Field Name	Description	Values/Range
EA	Specifies the extensible attribute	String
HWTYP	Specifies the hardware type	Example: IB-4015
MAX_DB_OBJECTS	Specifies the maximum objects in the database for a host	eg: 8000000
MAX_DHCP_LPS	Specifies the maximum number of DHCP leases per second for a host	Example: 15.0
MAX_DNS_QPS	Specifies the maximum DNS queries per second for a host	Example: 1000000.0
MEMBER_IP	Specifies the IP address of the member	IP address
timeendpos	Specifies the byte at which the timestamp ends. These values are based on the TIME_FORMAT that is specified for a sourcetype.	Example: 26
timestartpos	Specifies the byte at which the timestamp starts	Example: 0



## Indexes and Extracted Data

### Infoblox Audit Logs

Most of the fields in this index are extracted directly from the **audit.log** file. Some of them are mentioned in the table below:

Extracted Field Name	Description of the field	Values/Range	Source of Data
ACTION	Indicates the action taken	String. Example: Called	Infoblox audit logs
ADMIN	Indicates the name of the admin	String. Example: root	Infoblox audit logs
EA	Common Extracted Fields		
EXEC_STATUS	Indicates the execution status	String. Example: Pending Approval	Infoblox audit logs
HWTYP	Common Extracted Fields		
MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER_IP	Common Extracted Fields		
MESSAGE	Indicates the message	String. Example: to=Serial 040Console apparently_via=Direct auth=Local group=.admin-group	Infoblox audit logs
OBJECT_NAME	Indicates the object name	String. Example: RequestRestartServiceStatus	Infoblox audit logs
OBJECT_TYPE	Indicates the object type	String. Example: Shared AAAA Record	Infoblox audit logs
TIMESTAMP	Indicates the timestamp	Timestamp. Example: 2017-01-31 01:57:05	Infoblox audit logs
action	Indicates the action	Example: update, insert	Infoblox audit logs
address		Example: 10.0.0.0	Infoblox audit logs
auth		Example: Local	Infoblox audit logs
cidr		Example: 8	Infoblox audit logs

code		Example: created	Infoblox audit logs
comment		String	Infoblox audit logs
date_hour	Splunk Default field		
date_mday	Splunk Default field.		
date_minute	Splunk Default field		
date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
eventtype	Splunk Default field		
group		Example: admin-group	Infoblox audit logs
host	Splunk Default field		
index	Splunk Default field		
linecount	Splunk Default field		
member		Example: Member:infoblox.localdomain	Infoblox audit logs
network_view		Example: default	Infoblox audit logs
punct	Splunk Default field		
source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
user		Example: admin	Infoblox audit logs

Response_Time		Example: 0.1659	Infoblox audit logs
URI		Example: v2.10/record:host	Infoblox audit logs
InData		Example: {'comment': 'this is my one.xyz comment', 'name': 'user.zone.com', 'ipv4addrs': [{'configure_for_dhcp': False, 'mac': 'aa:0:0:0:1:cc', 'ipv4addr': '1.1.1.0'}], 'view': 'default'}	Infoblox audit logs

### Infoblox DNS Query, DNS Performance, DDNS, DNS Record Scavenging

Extracted Field Name	Description of the field	Values/Range	Source of Data
CLIENT	Indicates the DNS client	String	Infoblox DNS query
COUNT	Indicates the count	Integer	Infoblox DNS query and DNS Record Scavenging
EA	Common Extracted Fields		
FQDN	Indicates the FQDN	String	Infoblox DNS query
HITS	Indicates the DNS cache hits count	Integer	Infoblox DNS query
HNAME	Indicates the HNAME	String	Infoblox DNS query
HWTYP	Common Extracted Fields		
LATENCY	Indicates the latency count	Integer	Infoblox DNS performance
MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER	Specifies the member	String	DNS Record Scavenging
MEMBER_IP	Common Extracted fields		
MISSES	Specifies DNS cache miss count	Integer	Infoblox DNS query
QCOUNT	Specifies query count	Integer	Infoblox DNS query
REST	REST	String	Infoblox DDNS

SOURCE	SOURCE	String	Infoblox DDNS
SOURCEA	SOURCEA	IP address	Infoblox DDNS
TLD	Specifies the top-level domain name	String	Infoblox DNS query
TYPE	RR Type	String. Example: nxdomain	Infoblox DNS query and DNS Record Scavenging
TYPEA	TYPEA	String. Example: Success	Infoblox DDNS
VIEW		String	Infoblox DNS query
ZONE	Indicates the name of the zone	String	Infoblox DDNS
date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	Splunk Default field		
date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
display_name	Specifies the name of the DNS view	String	.
eventtype	Splunk Default field		
failure	Specifies the DNS FAILURE query count	Integer	
host	Splunk Default field		
index	Splunk Default field		
linecount	Splunk Default field		
nxdomain	Specifies the DNS NXDOMAIN query count	Integer	
nrrset	Specifies the DNS NXRRSET query count	Integer	

other	Specifies the DNS other query count	Integer	
punct	<i>Splunk Default field</i>		
referral	Specifies the DNS REFERRAL query count	Integer	
source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
success	Specifies the DNS success query count		
timeendpos	Common Extracted Fields		
timestartpos	Common Extracted Fields		

### Infoblox DNS Query Capture

Extracted Field Name	Description of the field	Values/Range	Source of Data
EA	Common Extracted Fields		
HWTYPE	Common Extracted Fields		
MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER_IP	Common Extracted Fields		
answer_count	Specifies the answer count	Integer	Infoblox DNS query capture
date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	Splunk Default field		
date_month	Splunk Default field		

date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
display_name	Specifies the DNS view	String	
eventtype	Splunk Default field		
flag_aa	Flag AA	Boolean. Example: Y	Infoblox DNS query capture
flag_ad	Flag AD	Boolean. Example: Y	Infoblox DNS query capture
flag_edns	Flag EDNS	Boolean. Example: Y	Infoblox DNS query capture
flag_recursion	Flag Recursion	Boolean. Example: Y	Infoblox DNS query capture
host	Splunk Default field		
host_class	Specifies the host class	Example: IN	Infoblox DNS query capture
host_type	Specifies the host type	Example: PTR	Infoblox DNS query capture
index	Splunk Default field		
linecount	Splunk Default field		
message_type	Specifies the message type	Example: Query or Response	Infoblox DNS query capture
name	Specifies the name	Host name. Example: 1.0.0.127.in-addr.arpa	Infoblox DNS query capture
query	Specifies the query	Host name. Example: 213.31.102.10.in-addr.arpa	Infoblox DNS query capture
query_class	Specifies the query class	Example: IN	Infoblox DNS query capture
query_count	Specifies the query count	Integer. Example: 1	Infoblox DNS query capture
query_source	Specifies the query source	Example: I, E	Infoblox DNS query capture
query_type	Specifies the DNS query type	Example: PTR	Infoblox DNS query capture

rdata	RDATA	String. This value depends on the query type.	Infoblox DNS query capture
reply_code	Specifies the reply code	String. Example: ServFail, NoError	Infoblox DNS query capture
source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
src_ip	Specifies the source IP	IP Address	Infoblox DNS query capture
src_port	Specifies the source port	Integer	Infoblox DNS query capture
time_msec	Specifies time in milliseconds	Integer	Infoblox DNS query capture
timeendpos	Common Extracted Fields		
timestamp	Indicates the timestamp	Integer	Infoblox DNS query capture
timestartpos	Common Extracted Fields		
transport	Specifies the mode of transport	Example: UDP, TCP	Infoblox DNS query capture
ttl	Specifies the TTL	Integer. Example: 3600	Infoblox DNS query capture
view	Specifies the view	Example: 1, 2	Infoblox DNS query capture

## Infoblox DHCP Performance

Extracted Field Name	Description of the field	Values/Range	Source of Data
EA	Common Extracted Fields		
HWTYP	Common Extracted Fields		
MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		

MEMBER_IP	Common Extracted Fields		
NETWORK	Specifies the network address	Example: 10.0.0.0/8	
address	Specifies the DHCP client address	IP address	Infoblox DHCP performance
address_total	Specifies the total number of addresses	Integer	Infoblox DHCP performance
cidr	Specifies the CIDR	Example: 24	Infoblox DHCP performance
date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	Splunk Default field		
date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
dhcp_hosts	Specifies the DHCP hosts count	Integer	Infoblox DHCP performance
dhcp_utilization	Specifies the DHCP utilization	Integer	Infoblox DHCP performance
dhcp_utilization_status	Specifies the DHCP utilization status	String	Infoblox DHCP performance
dhcpv4ack	Specifies the DHCPv4 ACK message count	Integer	Infoblox DHCP performance
dhcpv4decline	Specifies the DHCPv4 decline message count	Integer	Infoblox DHCP performance
dhcpv4discover	Specifies the DHCPv4 discover message count	Integer	Infoblox DHCP performance
dhcpv4inform	Specifies the DHCPv4 inform message count	Integer	Infoblox DHCP performance
dhcpv4leaseactive	Specifies the DHCPv4 lease active message count	Integer	Infoblox DHCP performance



dhcpv4leasequery	Specifies the DHCPv4 lease query message count	Integer	Infoblox DHCP performance
dhcpv4leaseunassigned	Specifies the DHCPv4 lease unassigned message count	Integer	Infoblox DHCP performance
dhcpv4leaseunknown	Specifies the DHCPv4 lease unknown message count	Integer	Infoblox DHCP performance
dhcpv4nak	Specifies the DHCPv4 NAK message count	Integer	Infoblox DHCP performance
dhcpv4offer	Specifies the DHCPv4 offer message count	Integer	Infoblox DHCP performance
dhcpv4release	Specifies the DHCPv4 release message count	Integer	Infoblox DHCP performance
dhcpv4request	Specifies the DHCPv4 request message count	Integer	Infoblox DHCP performance
dhcpv6advertise	Specifies the DHCPv6 advertise message count	Integer	Infoblox DHCP performance
dhcpv6confirm	Specifies the DHCPv6 confirm message count	Integer	Infoblox DHCP performance
dhcpv6decline	Specifies the DHCPv6 decline message count	Integer	Infoblox DHCP performance
dhcpv6information_request	Specifies the DHCPv6 information request message count	Integer	Infoblox DHCP performance
dhcpv6leasequery	Specifies the DHCPv6 lease query message count	Integer	Infoblox DHCP performance
dhcpv6leasequery_reply	Specifies the DHCPv6 lease query reply message count	Integer	Infoblox DHCP performance
dhcpv6rebind	Specifies the DHCPv6 rebind message count	Integer	Infoblox DHCP performance
dhcpv6reconfigure	Specifies the DHCPv6 reconfigure message count	Integer	Infoblox DHCP performance
dhcpv6relay_forward	Specifies the DHCPv6 relay forward message count	Integer	Infoblox DHCP performance
dhcpv6relay_reply	Specifies the DHCPv6 relay reply message count	Integer	Infoblox DHCP performance
dhcpv6release	Specifies the DHCPv6 release message count	Integer	Infoblox DHCP performance

dhcpv6renew	Specifies the DHCPv6 renew message count	Integer	Infoblox DHCP performance
dhcpv6reply	Specifies the DHCPv6 reply message count	Integer	Infoblox DHCP performance
dhcpv6request	Specifies the DHCPv6 request message count	Integer	Infoblox DHCP performance
dhcpv6solicit	Specifies the DHCPv6 solicit message count	Integer	Infoblox DHCP performance
display_name	Specifies the DNS View	String	
dynamic_hosts	Specifies the dynamic hosts count	Integer	Infoblox DHCP performance
end_address	Specifies the end IP address	IP address	Infoblox DHCP performance
eventtype	Splunk Default field		
host	Splunk Default field		
index	Splunk Default field		
linecount	Splunk Default field		
members	Specifies the DHCP member	Example: infoblox.localdomain	Infoblox DHCP performance
ms_servers	Specifies the MS servers	IP address	Infoblox DHCP performance
protocol	Specifies the DHCP protocol	Example: IPV4	
punct	Splunk Default field		
ranges	Specifies the DHCP ranges count	Integer	Infoblox DHCP performance
source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
start_address	Specifies the start IP address	IP address	Infoblox DHCP performance
static_hosts	Specifies the static hosts count	Integer	Infoblox DHCP performance

timeendpos	Common Extracted Fields		
timestamp	Specifies the timestamp of the event	Example: 2017-02-04 03:45:53	Infoblox DHCP performance
timestartpos	Common Extracted Fields		
View	Specifies the network view	Example: default	Infoblox DHCP performance

### Infoblox DHCP FingerPrint, DHCP Lease History

Extracted Field Name	Description of the field	Values/Range	Source of Data
ACTION	Specifies the action	String. Example: Issued	Infoblox DHCP lease history
CIDR	Specifies the CIDR	Integer	Infoblox DHCP lease history
DEVICE_CLASS	Specifies the device class	String. Example: Linux	
EA	Common Extracted Fields		
END_EPOCH	Specifies the end epoch time	Integer	Infoblox DHCP lease history
FP	Specifies the name of the DHCP fingerprint	String. Example: No Match	Infoblox DHCP lease history
FP_CIDR	Specifies the fingerprint CIDR	Integer. Example: 8	Infoblox DHCP lease history
FP_NW	Specifies the fingerprint network	Network address. Example: 10.0.0.0	Infoblox DHCP lease history
FP_RANGE	Specifies the fingerprint range	Network range. Example: 10.0.0.1-10.0.0.200	Infoblox DHCP lease history
FP_VIEW	Specifies the fingerprint view	String. Example: default	Infoblox DHCP lease history
HWTYPE	Common Extracted fields		
LEASE_IP	Specifies the lease IP address	IP address	Infoblox DHCP lease history
MAC_DUID	Specifies the MAC address	MAC address	Infoblox DHCP lease history
MAX_DB_OBJECTS	Common Extracted fields		
MAX_DHCP_LPS	Common Extracted fields		

MAX_DNS_QPS	Common Extracted fields		
MEMBER_IP	Common Extracted fields		
MS Server	Specifies the MS server	IP Address	Infoblox DHCP lease history
NW	Specifies the network	Network address. Example: 10.0.0.0	Infoblox DHCP lease history
OPTION12HOST	Specifies the host name that is sent using DHCP Option 12	String. Example: Fedora21	Infoblox DHCP lease history
OS_NUMBER	Specifies the OS number	Integer	Infoblox DHCP lease history
PROTO	Specifies the protocol	String. Example: dhcpd	Infoblox DHCP lease history
SFP	SFP	String. Example: Ubuntu/Debian 5/ Knoppix 6	Infoblox DHCP fingerprint
START_EPOCH	Specifies the start epoch time	Integer	Infoblox DHCP lease history
VIEW	Specifies the view		Infoblox DHCP lease history
date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	Splunk Default field		
date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
display_name	Specifies the DNS view	String	
eventtype	Splunk Default field		
host	Splunk Default field		
index	Splunk Default field		

linecount	Splunk Default field		
punct	Splunk Default field		
source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
timeendpos	Common extracted fields		
timestartpos	Common extracted fields		

### Infoblox DDI Utilization

Extracted Field Name	Description of the field	Values/Range	Source of Data
EA	Common Extracted Fields		
HWTYP	Common Extracted Fields		
MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER_IP	Common Extracted Fields		
address_alloc	Specifies the address allocation count	Integer	Infoblox DDI utilization
address_assignable	Specifies the address assignable count	Integer	Infoblox DDI utilization
address_assigned	Specifies the address assigned count	Integer	Infoblox DDI utilization
address_conflicts	Specifies the address conflicts count		Infoblox DDI utilization
address_reserved	Specifies the address reserved count	Integer	Infoblox DDI utilization
address_total	Specifies the total number of addresses	Integer	Infoblox DDI utilization

address_unalloc	Specifies the address unallocation count	Integer	Infoblox DDI utilization
address_unmanaged	Specifies the address unmanaged count	Integer	Infoblox DDI utilization
allocation	Allocation	Integer	Infoblox DDI utilization
cidr	Specifies the CIDR	Example: 24	Infoblox DDI utilization
date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	Splunk Default field		
date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
discovered_name	Specifies the discovered name	String	Infoblox DDI utilization
display_name	Specifies the DNS view	String	
eventtype	Splunk Default field		
first_discovered_timestamp	Specifies the first discovered timestamp	Timestamp	Infoblox DDI utilization
host	Splunk Default field		
hosts	Specifies the address hosts count	Integer	Infoblox DDI utilization
index	Splunk Default field		
ip_address	Specifies the IP address	IP Address	Infoblox DDI utilization
last_discovered_timestamp	Specifies the last discovered timestamp	timestamp	Infoblox DDI utilization
linecount	Splunk Default field		

managed	Indicates if managed or not	Boolean	Infoblox DDI utilization
management_platform	Specifies the management platform	String	Infoblox DDI utilization
members	Specifies the DHCP members	Example: infoblox.localdomain	Infoblox DDI utilization
ms_primary	Specifies the MS primary	String	Infoblox DDI utilization
port_vlan_name	Specifies the VLAN port name	String	Infoblox DDI utilization
port_vlan_number	Specifies the VLAN port number	Integer	Infoblox DDI utilization
network_view	Specifies the network view	String	Infoblox DDI utilization
primary	Primary	FQDN	Infoblox DDI utilization
protocol	Specifies the DHCP protocol	Example: IPV4	Infoblox DDI utilization
punct	Splunk Default field		
rr_a	Specifies the resource record A count	Integer	Infoblox DDI utilization
rr_aaaa	Specifies the resource record AAAA count	Integer	Infoblox DDI utilization
rr_cname	Specifies the resource record CNAME count	Integer	Infoblox DDI utilization
rr_dhcid	Specifies the resource record DHCID count	Integer	Infoblox DDI utilization
rr_dname	Specifies the resource record DNAME count	Integer	Infoblox DDI utilization
rr_dnskey	Specifies the resource record DNSKEY count	Integer	Infoblox DDI utilization
rr_ds	Specifies the resource record DS count	Integer	Infoblox DDI utilization
rr_lbdn	Specifies the resource record LBDN count	Integer	Infoblox DDI utilization
rr_mx	Specifies the resource record MX count	Integer	Infoblox DDI utilization
rr_naptr	Specifies the resource record NAPTR count	Integer	Infoblox DDI utilization
rr_ns	Specifies the resource record NS count	Integer	Infoblox DDI utilization

rr_nsec	Specifies the resource record NSEC count	Integer	Infoblox DDI utilization
rr_nsec3	Specifies the resource record NSEC3 count	Integer	Infoblox DDI utilization
rr_nsec3param	Specifies the resource record NSEC3PARAM count	Integer	Infoblox DDI utilization
rr_other	Specifies the resource record OTHER count	Integer	Infoblox DDI utilization
rr_ptr	Specifies the resource record PTR count	Integer	Infoblox DDI utilization
rr_rrsig	Specifies the resource record RRSIG count	Integer	Infoblox DDI utilization
rr_soa	Specifies the resource record SOA count	Integer	Infoblox DDI utilization
rr_srv	Specifies the resource record SRV count	Integer	Infoblox DDI utilization
rr_tlsa	Specifies the resource record TLSA count	Integer	Infoblox DDI utilization
rr_total	Specifies the resource record TOTAL count	Integer	Infoblox DDI utilization
rr_txt	Specifies the resource record TXT count	Integer	Infoblox DDI utilization
signed	Indicates whether signed or not	Boolean	Infoblox DDI utilization
source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
timeendpos	Common Extracted Fields		
Timestamp	Specifies the timestamp of the event	Example: 2017-02-04 03:45:53	Infoblox DDI utilization
timestartpos	Common Extracted Fields		
utilization	Specifies the address utilization count	Integer	Infoblox DDI utilization



view	Specifies the network view	Example: default	Infoblox DDI utilization
zone_format	Specifies the zone format	String. Example: Forward-Mapping	Infoblox DDI utilization
zone_name	Specifies the zone name	String. Example: member1.com	Infoblox DDI utilization
zones_forward	Specifies the zone forward count	Integer	Infoblox DDI utilization
zones_ipv4	Specifies the IPv4 count of the zone	Integer	Infoblox DDI utilization
zones_ipv6	Specifies the IPv6 count of the zone	Integer	Infoblox DDI utilization
zones_signed	Specifies the signed count of the zone	Integer	Infoblox DDI utilization

### Infoblox Discovered Devices Related Dashboards/Reports

Extracted Field Name	Description of the field	Values/Range	Source of Data
ADM_DN_OP_DN_COUNT	Admin-Down/Operation-DownPort Count	Integer	Infoblox discovered devices related dashboards/reports
ADM_UP_OP_DN_COUNT	Admin-Up/Operation-UpPort Count	Integer	Infoblox discovered devices related dashboards/reports
ADM_UP_OP_UP_COUNT	Admin-Up/Operation-DownPort Count	Integer	Infoblox discovered devices related dashboards/reports
COMPONENT_NAME	Specifies the component name	String. Example: DELL-PC8024F	Infoblox discovered devices related dashboards/reports
COMPONENT_TYPE	Specifies the component type	String. Example: Switch-Router	Infoblox discovered devices related dashboards/reports
COMPONENT_PORT	Specifies the component port	String. Example: Gi1/0/24	Infoblox discovered devices related dashboards/reports
DEVICE_MGMT_IP	Specifies the device management IP address	IP address	Infoblox discovered devices related dashboards/reports
DEVICE_MODEL	Specifies the device model	String. Example: EX2200	Infoblox discovered devices related dashboards/reports

DEVICE_NAME	Specifies the device name	String. Example: Cisco_434f44	Infoblox discovered devices related dashboards/reports
DEVICE_TYPE	Specifies the device type	String. Example: Switch, Router	Infoblox discovered devices related dashboards/reports
DEVICE_VENDOR	Specifies the device vendor	String. Example: Avaya	Infoblox discovered devices related dashboards/reports
DISCOVERED_MAC_DUID	Specifies the discovered MAC DUID	MAC address	Infoblox discovered devices related dashboards/reports
DISCOVERED_NAME	Specifies the discovered name	Example: dev_view1.yahoo.com	Infoblox discovered devices related dashboards/reports
EA	Common Extracted Fields		
HWTYP	Common Extracted Fields		
IN_USE_FLAG	In use flag	Integer. Example: 1	Infoblox discovered devices related dashboards/reports
IPADDR	Specifies the IP address	IP Address. Example: 11.11.11.11	Infoblox discovered devices related dashboards/reports
IPADDR_MASK	Specifies the IP address mask	Integer. Example: 128	Infoblox discovered devices related dashboards/reports
MAC_DUID	Specifies the MAC address	MAC address	Infoblox discovered devices related dashboards/reports
MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER_IP	Common Extracted Fields		
NETWORK_VIEW	Specifies the network view	String. Example: default	Infoblox discovered devices related dashboards/reports
NON_NULL_NAME	Specifies the non-null name	String. Example: DELL-PC8024F	

NON_NULL_PORT	Specifies the non-null port	String. Example: Gi1/0/24	
TIMESTAMP	Specifies the timestamp	Timestamp. Example: 2017-02-15 15:56:27	Infoblox discovered devices related dashboards/reports
TIMESTAMP_USER_HOST_PROCESS_PID_INFO_PREFIX	Specifies the timestamp userhost process pid info prefix	String. Example: 2017-02-15T11:02:53+00:00 user infoblox.localdomain <ac:structured-macro ac.name="unmigrated-wiki-markup" ac.schema-version="1" ac:macro-id="e5d51f7e-f354-4235-870a-9e02f49b3d41"><ac:plain-text-body><![CDATA[python[]: info ipaddr-activity-rpt	Infoblox discovered devices related dashboards/reports
TOTAL_AVAIL_COUNT	Specifies the total available count	Integer	Infoblox discovered devices related dashboards/reports
Type	Specifies the type	String. Example: Discovery	
ap_bss_mac	Access Point BSS MAC	MAC address	Infoblox discovered devices related dashboards/reports
ap_ip_dotted	Access Point IP dotted	String	Infoblox discovered devices related dashboards/reports
ap_mac	Access Point MAC	MAC address	Infoblox discovered devices related dashboards/reports
ap_name	Access Point name	String	Infoblox discovered devices related dashboards/reports
ap_associated_ssid	Access Point associated SSID	String	Infoblox discovered devices related dashboards/reports
asset_type	Specifies the asset type	String. Example: Physical Device	Infoblox discovered devices related dashboards/reports
class	Specifies the class name	String. Example: port	Infoblox discovered devices related dashboards/reports
component_name	Specifies the component name	String. Example: GigabitEthernet1/0/1	Infoblox discovered devices related dashboards/reports
date_hour	Splunk Default field		

date_mday	Splunk Default field		
date_minute	Splunk Default field		
date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
Description	Specifies the description	String. Example: Gigabit Ethernet Port	Infoblox discovered devices related dashboards/reports
device_id	Specifies the device ID	Integer	Infoblox discovered devices related dashboards/reports
device_ip_address	Specifies the device IP address	IP address	Infoblox discovered devices related dashboards/reports
device_model	Specifies the device model	String. Example: catalyst37xxStack	Infoblox discovered devices related dashboards/reports
device_name	Specifies the device name	String. Example:DELL-PC8024F	Infoblox discovered devices related dashboards/reports
device_os_version	Specifies the device OS version	String. Example: 4.14.6M	Infoblox discovered devices related dashboards/reports
device_type	Specifies the device type	String. Example: Switch	Infoblox discovered devices related dashboards/reports
device_vendor	Specifies the device vendor	String. Example: Cisco	Infoblox discovered devices related dashboards/reports
device_version	Specifies the device version	String. Example: 5.1.2.3	Infoblox discovered devices related dashboards/reports
display_name	Specifies the DNS view	String	Infoblox discovered devices related dashboards/reports

end_host_addl_info	Specifies additional information about the end host	String	Infoblox discovered devices related dashboards/reports
end_host_device_model	Specifies the device model of the end host	String. Example: catalyst37xxStack	Infoblox discovered devices related dashboards/reports
end_host_device_type	Specifies the device type of the end host	String. Example: Switch-Router	Infoblox discovered devices related dashboards/reports
end_host_device_vendor	Specifies the device vendor of the end host	String. Example: Cisco	Infoblox discovered devices related dashboards/reports
end_host_first_discovered	Specifies the first occasion when the end host was first discovered	Integer	Infoblox discovered devices related dashboards/reports
end_host_ip_address	Specifies the IP address of the end host	IP address	Infoblox discovered devices related dashboards/reports
end_host_last_discovered	Indicates when was end host last discovered	Integer	Infoblox discovered devices related dashboards/reports
end_host_mac_address	Specifies the MAC address of the end host	MAC address	Infoblox discovered devices related dashboards/reports
end_host_name	Specifies the name of the end host	String. Example: WS-C3750X-24P	Infoblox discovered devices related dashboards/reports
end_host_network_view	Specifies the network view of the end host	String. Example: custom view	Infoblox discovered devices related dashboards/reports
end_host_os_version	Specifies the version of the end host OS	String. Example: 15.2(1)E2	Infoblox discovered devices related dashboards/reports
eventtype	Splunk Default field		
firmware_rev	Indicates firmware revision	String. Example: 15.2(1)E2	Infoblox discovered devices related dashboards/reports
first_seen	First seen timestamp	Integer	Infoblox discovered devices related dashboards/reports
hardware_rev	Specifies revision of the hardware	String. Example: V05	Infoblox discovered devices related dashboards/reports
host	Splunk Default field		
index	Splunk Default field		
interface_admin_status	Specifies the interface admin status	String. Example: up	Infoblox discovered devices related dashboards/reports

interface_description	Specifies the interface interface description	String	Infoblox discovered devices related dashboards/reports
interface_ip_address	Specifies the interface IP address	IP address	Infoblox discovered devices related dashboards/reports
interface_name	Specifies the interface name	String. Example: Fa0	Infoblox discovered devices related dashboards/reports
interface_port_status	Specifies the interface port status	String. Example: up	Infoblox discovered devices related dashboards/reports
interface_speed	Specifies the interface speed	Integer. Example: 1000000000	Infoblox discovered devices related dashboards/reports
interface_type	Specifies the interface type	String. Example: tunnel	Infoblox discovered devices related dashboards/reports
interface_vlan	Specifies the interface VLAN ID	Integer Example: 16	Infoblox discovered devices related dashboards/reports
interface_vlan_name	Specifies the interface VLAN name	String. Example: VLAN1014	Infoblox discovered devices related dashboards/reports
ip_address	Specifies the IP address	IP address	Infoblox discovered devices related dashboards/reports
is_trunk_port	Specifies if it is a trunk port or not	Boolean	Infoblox discovered devices related dashboards/reports
last_seen	Specifies the last seen timestamp	Integer	Infoblox discovered devices related dashboards/reports
linecount	Splunk Default field		
model	Specifies the model name	String. Example: DCS-7048T-A	Infoblox discovered devices related dashboards/reports
network_view	Specifies the network view	String. Example: custom view	Infoblox discovered devices related dashboards/reports
port_last_changed_at	The timestamp when the port was last changed	Timestamp	Infoblox discovered devices related dashboards/reports
punct	Splunk Default field		
serial_number	Specifies the serial number	String. Example: JPE12440180	Infoblox discovered devices related dashboards/reports
software_rev	Specifies the software revision	String. Example: 15.2(1)E2	Infoblox discovered devices related dashboards/reports

source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
switch_interface	Specifies the switch interface	String. Example: Gi0/47	Infoblox discovered devices related dashboards/reports
switch_ip_address	Specifies the switch IP Address	IP Address	Infoblox discovered devices related dashboards/reports
switch_model	Indicates the switch model	String. Example: cat3560x48	Infoblox discovered devices related dashboards/reports
switch_name	Specifies the switch name	String. Example: ni-mri-sw4.inca.infoblox.com	Infoblox discovered devices related dashboards/reports
switch_os_version	Specifies the OS version of the switch	String. Example: 12.2(53)SE2	Infoblox discovered devices related dashboards/reports
switch_type	Specifies the switch type	String. Example: Switch	Infoblox discovered devices related dashboards/reports
switch_vendor	Specifies the vendor of the switch	String. Example: Cisco	Infoblox discovered devices related dashboards/reports
switch_vlan	Specifies the switch VLAN	Integer. Example: 18	Infoblox discovered devices related dashboards/reports
timeendpos	Common Extracted Fields		
timestamp	Indicates the timestamp	Integer	Infoblox discovered devices related dashboards/reports
timestamp_user_host_process_pid_info_prefix	Specifies the prefix	String	Infoblox discovered devices related dashboards/reports
timestartpos	Common Extracted Fields		
user_id	Specifies the User ID		Infoblox discovered devices related dashboards/reports
View	Specifies the DNS view	String	Infoblox discovered devices related dashboards/reports
virtual_ind	Specifies the virtual indicator	Integer	

## Infoblox Threat Protection Related Dashboards/Reports

Extracted Field Name	Description of the field	Values/Range	Source of Data
ACOUNT	ACOUNT	Integer	Infoblox threat protection related dashboards/reports
ACTIVE_COUNT	Specifies the active count	Integer	Infoblox threat protection related dashboards/reports
ALERT_ID	Specifies the alert ID	Integer	Infoblox threat protection related dashboards/reports
ALERT_TYPE	Specifies the alert type	String	Infoblox threat protection related dashboards/reports
BLOCK_END	Specifies the block end IP address	Integer	Infoblox threat protection related dashboards/reports
BLOCK_START	Specifies the block start IP address	Integer	Infoblox threat protection related dashboards/reports
CATEGORY	Specifies the category	String. Example: OSPF	Infoblox threat protection related dashboards/reports
CLIENT	Specifies the client	String	Infoblox threat protection related dashboards/reports
COUNT	Specifies the count	Integer	Infoblox threat protection related dashboards/reports
DCOUNT	Specifies the DCOUNT	Integer	Infoblox threat protection related dashboards/reports
DNST_CATEGORY	Specifies the destination category	String	Infoblox threat protection related dashboards/reports
DOMAIN_NAME	Specifies the domain name	String	Infoblox threat protection related dashboards/reports
EA	Common Extracted Fields		
FIREEYE_APPLIANCE	Specifies the FireEye appliance	String	Infoblox threat protection related dashboards/reports
HWTYPE	Common Extracted Fields		
LOG_SEVERITY	Specifies log severity	String	Infoblox threat protection related dashboards/reports
MAX_DB_OBJECTS	Common Extracted Fields		



MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER_IP	Common Extracted Fields		
MESSAGE	Specifies the message	String. Example: DROP OSPF unexpected	Infoblox threat protection related dashboards/reports
MITIGATION_ACTION	Specifies the mitigation action	String	Infoblox threat protection related dashboards/reports
NAT_STATUS	Specifies the NAT status	String	Infoblox threat protection related dashboards/reports
RECORD_DATA	Specifies the record data	String	Infoblox threat protection related dashboards/reports
RPZ_QNAME	Specifies the RPZ QNAME	String	Infoblox threat protection related dashboards/reports
RULE_DESCRIPTION	Specifies the rule description	String. Example: This rule drops any unexpected OSPF packets when OSPF is disabled.	
RULE_NAME	Specifies the rule name	String. Example: DROP OSPF unexpected	
RULE_SID	Specifies the rule SID	Integer	Infoblox threat protection related dashboards/reports
SEVERITY	Specifies the severity	String. Example: INFORMATIONAL	Infoblox threat protection related dashboards/reports
SID	Specifies the SID	Integer	Infoblox threat protection related dashboards/reports
SOURCE_IP	Specifies the source IP	IP address	Infoblox threat protection related dashboards/reports
SOURCE_PORT	Specifies the source port	Integer	Infoblox threat protection related dashboards/reports
TIMESTAMP	Indicates the timestamp	Timestamp	Infoblox threat protection related dashboards/reports
TOTAL_COUNT	Specifies the total count	Integer	Infoblox threat protection related dashboards/reports
VIEW	Specifies the DNS view	String	Infoblox threat protection related dashboards/reports

date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	Splunk Default field		
date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
eventtype	Splunk Default field		
host	Splunk Default field		
index	Splunk Default field		
linecount	Splunk Default field		
punct	Splunk Default field		
source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
timestartpos	Common extracted fields		

### Infoblox DNS Traffic Control

Most of the fields in this index are extracted directly from the **syslog\_filtered.log** file. Some of them are mentioned in the table below:

Extracted Field Name	Description of the field	Values/Range	Source of Data
EA	Common Extracted Fields		
HWTYP	Common Extracted Fields		

MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER_IP	Common Extracted Fields		
available	<i>Specifies the available count</i>	Integer	Infoblox DNS traffic control
date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	Splunk Default field		
date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
eventtype	Splunk Default field		
host	Splunk Default field		
index	Splunk Default field		
linecount	Splunk Default field		
monitor	Specifies the DNS Traffic Control SNMP health monitor	String	Infoblox DNS traffic control
pool	Specifies the pool	String	Infoblox DNS traffic control
punct	Splunk Default field		
resource	Specifies the resource	String	Infoblox DNS traffic control
response_count	Specifies the response count	Integer	Infoblox DNS traffic control
source	Splunk Default field		

sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
timeendpos	Common Extracted Fields		
timestamp	Indicates the timestamp of the event	Example: 2017-02-04 03:45:53	
timestartpos	Common Extracted Fields		
unavailable	Specifies the unavailable count	Integer	Infoblox DNS traffic control

### Infoblox Cloud Related Dashboards/Reports

Extracted Field Name	Description of the field	Values/Range	Source of Data
ACTION	Specifies the action	String. Example: Allocated	
EA	Common Extracted Fields		
HWTYPE	Common Extracted Fields		
MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER_IP	Common Extracted Fields		
TENANT_NAME	Specifies the name of the tenant associated with the VM	String	
action	Specifies the action count	Integer	Infoblox cloud related dashboards/ reports
address	Specifies the IP address	IP address	Infoblox cloud related dashboards/ reports
address_type	Specifies the type of address	Integer	Infoblox cloud related dashboards/ reports
application_type	Specifies the application type		Infoblox cloud related dashboards/ reports

cidr	Specifies the CIDR	Example: 24	Infoblox cloud related dashboards/ reports
cnames	Specifies the common name	String	Infoblox cloud related dashboards/ reports
date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	Splunk Default field		
date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
display_name	Specifies the DNS view	String	
elastic_address	Specifies the elastic IP address	IP address	Infoblox cloud related dashboards/ reports
eventtype	Splunk Default field		
Fqdn	Specifies the FQDN	String	Infoblox cloud related dashboards/ reports
host	Splunk Default field		
index	Splunk Default field		
interface_name	Specifies the interface name	String	Infoblox cloud related dashboards/ reports
is_primary_ifc	Indicates if primary IFC or not	Example: 0 (not primary)	Infoblox cloud related dashboards/ reports
linecount	Splunk Default field		
location	Specifies the location		Infoblox cloud related dashboards/ reports
mac_address	Specifies the MAC address	Example: 00:11:22:33:44:55	Infoblox cloud related dashboards/ reports

mgmt_platform	Specifies management platform	Example: vm132ctest	Infoblox cloud related dashboards/ reports
network	Specifies the network address	Example: 10.0.0.0/8	Infoblox cloud related dashboards/ reports
network_view	Specifies the network view	Example: default	Infoblox cloud related dashboards/ reports
port_id	Specifies the port ID	Integer	Infoblox cloud related dashboards/ reports
private_address	Specifies the private address	IP address	Infoblox cloud related dashboards/ reports
private_hostname	Specifies the private hostname	String	Infoblox cloud related dashboards/ reports
public_address	Specifies the public address	IP address	Infoblox cloud related dashboards/ reports
public_hostname	Specifies the public hostname	String	Infoblox cloud related dashboards/ reports
punct	Splunk Default field		
source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
tenant_id	Specifies the tenant ID	Integer	Infoblox cloud related dashboards/ reports
timeendpos	Common Extracted Fields		
timestamp	Indicates the timestamp of the event	Example: 2017-02-04 03:45:53	Infoblox cloud related dashboards/ reports
timestartpos	Common Extracted Fields		
view	Specifies the DNS view	String	
vlan_id	Specifies the VLAN ID	Integer	Infoblox cloud related dashboards/ reports

vm_hostname	Specifies the hostname of the VM	String	Infoblox cloud related dashboards/ reports
vm_name	Specifies the name of the VM	Example: 99	Infoblox cloud related dashboards/ reports
vm_vpc_address	Specifies the VPC address of the VM	IP address	Infoblox cloud related dashboards/ reports
vm_vpc_cidr	Specifies the VPC CIDR of the VM	Example: 24	Infoblox cloud related dashboards/ reports
vm_vpc_id	Specifies the VPC ID of the VM	Integer	Infoblox cloud related dashboards/ reports
vm_vpc_name	Specifies the VPC name of the VM	Integer	Infoblox cloud related dashboards/ reports
vpc_addr	Specifies the VPC address	IP address	Infoblox cloud related dashboards/ reports

## Infoblox Syslog

Most of the fields in this index are extracted directly from the **syslog\_filtered.log** file. Some of them are mentioned in the table below:

Extracted Field Name	Description of the field	Values/Range	Source of Data
BOOT_IMAGE		Example: /boot/bzImage	Infoblox syslog file
CPUs		Integer. Example: 8	Infoblox syslog file
EA	Common Extracted Fields		
HWTYPE	Common Extracted Fields		
MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER_IP	Common Extracted Fields		
date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	Splunk Default field		

date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
eventtype	Splunk Default field		
group		Example: admin-group	Infoblox syslog file
hits		Integer	Infoblox syslog file
host	Splunk Default field		
index	Splunk Default field		
linecount	Splunk Default field		
misses		Integer	Infoblox syslog file
punct	Splunk Default field		
size		Integer	Infoblox syslog file
source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
timeendpos	Common Extracted Fields		
timestartpos	Common Extracted Fields		

## System Capacity

Extracted Field Name	Description of the field	Values/Range	Source of Data
COUNT	Specifies the count	Integer	System capacity



EA	Common Extracted Fields		
HWTYPE	Common Extracted Fields		
MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER_IP	Common Extracted Fields		
PERCENT	Specifies the percentage	Integer	System capacity
date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	<i>Splunk Default field</i>		
date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
eventtype	Splunk Default field		
host	Splunk Default field		
index	Splunk Default field		
linecount	Splunk Default field		
punct	Splunk Default field		
source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		

splunk_server_group	Splunk Default field		
timeendpos	Common Extracted Fields		
timestartpos	Common Extracted Fields		

### Infoblox System Utilization (CPU, Memory, Network Traffic) Related Dashboards/Reports

Extracted Field Name	Description of the field	Values/Range	Source of Data
CPU_PERCENT	Specifies the CPU percentage	Integer value within 0-100	Infoblox system utilization related dashboards/reports
EA	Common Extracted Fields		
HWTYPE	Common Extracted Fields		
MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER_IP	Common Extracted Fields		
MEMORY_PERCENT	Specifies the memory percentage	Integer. Value within 0-100	Infoblox system utilization related dashboards/reports
TRAF_VALUE	Specifies the traffic value	Integer	Infoblox system utilization related dashboards/reports
date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	Splunk Default field		
date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		

eventtype	Splunk Default field		
host	Splunk Default field		
index	Splunk Default field		
linecount	Splunk Default field		
punct	Splunk Default field		
source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
sys_report_id	Specifies the report ID based on whether inbound or outbound	Integer	Infoblox system utilization related dashboards/reports
timeendpos	Common Extracted Fields		
timestartpos	Common Extracted Fields		

### Infoblox Ecosystem Subscription

Extracted Field Name	Description of the field	Values/Range	Source of Data
EA	Common Extracted Fields		
HWTYPE	Common Extracted Fields		
MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER_IP	Common Extracted Fields		
cisco_ise_endpoint_profile	Specifies the Cisco ISE endpoint profile	String	Infoblox ecosystem subscription
cisco_ise_security_group	Specifies the Cisco ISE security group		Infoblox ecosystem subscription

cisco_ise_session_state	Specifies the Cisco ISE session state	String. Example: STARTED	Infoblox ecosystem subscription
cisco_ise_ssid	Specifies the Cisco ISE SSID	String	Infoblox ecosystem subscription
date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	Splunk Default field		
date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
domainname	Specifies the domain name	String	Infoblox ecosystem subscription
ea_eps_status	Specifies the EPS status of the extensible attribute	String	Infoblox ecosystem subscription
eventtype	Splunk Default field		
guid	Specifies the GUID	String	Infoblox ecosystem subscription
host	Splunk Default field		
index	Splunk Default field		
ip_address	Specifies the IP address	IP address	Infoblox ecosystem subscription
last_discovered_timestamp	Specifies the last discovered timestamp	Integer	Infoblox ecosystem subscription
linecount	Splunk Default field		
port_vlan_name	Specifies the VLAN name of the port	String	Infoblox ecosystem subscription
port_vlan_number	Specifies the VLAN number of the port	Integer	Infoblox ecosystem subscription
punct	Splunk Default field		

source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
timeendpos	Common Extracted Fields		
timestamp	Specifies the timestamp of the event	Example: 2017-02-04 03:45:53	Infoblox ecosystem subscription
timestartpos	Common Extracted Fields		
username	Specifies the username	String	Infoblox ecosystem subscription

### Infoblox Ecosystem Publication

Extracted Field Name	Description of the field	Values/Range	Source of Data
EA	Common Extracted Fields		
HWTYP	Common Extracted Fields		
MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER_IP	Common Extracted Fields		
contents	Specifies the content	String. Example: {'LEASE_STATE': 'STARTED', 'Lease_Start_Time': '2017-03-01T07:00:00Z', 'MAC_OR_DUID': '80:3c:3e:29:84:cc', 'Fingerprint': 'No Match', 'Lease_End_Time': '2017-03-01T07:02:00Z', 'IPAddress': '10.0.0.20', 'Infoblox_Member': '10.35.205.6'}	Infoblox ecosystem publication
date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	Splunk Default field		

date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
eventtype	Splunk Default field		
host	Splunk Default field		
index	Splunk Default field		
ip_address	Specifies the IP address	IP address	Infoblox ecosystem publication
linecount	Splunk Default field		
notification_action	Specifies the notification action	Example: CISCOISE_PUBLISH_IPAM	Infoblox ecosystem publication
notification_target	Specifies the notification target	IP address	Infoblox ecosystem publication
punct	Splunk Default field		
source	Splunk Default field		
sourcetype	Splunk Default field		
splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
timeendpos	Common Extracted Fields		
timestamp	Specifies the timestamp of the event	Example: 2017-02-04 03:45:53	Infoblox ecosystem publication
timestartpos	Common Extracted Fields		

## Reporting License Usage

Extracted Field Name	Description of the field	Values/Range	Source of Data
EA	Common Extracted Fields		

HWTYP	Common Extracted Fields		
MAX_DB_OBJECTS	Common Extracted Fields		
MAX_DHCP_LPS	Common Extracted Fields		
MAX_DNS_QPS	Common Extracted Fields		
MEMBER_IP	Common Extracted Fields		
date_hour	Splunk Default field		
date_mday	Splunk Default field		
date_minute	Splunk Default field		
date_month	Splunk Default field		
date_second	Splunk Default field		
date_wday	Splunk Default field		
date_year	Splunk Default field		
date_zone	Splunk Default field		
display_name	Specifies the DNS view	String	
eventtype	Splunk Default field		
host	Splunk Default field		
index	Splunk Default field		
license_count	Specifies the license count	Integer	Reporting license usage
license_pool	Specifies the license pool	String. Example: cloud_api.0	Reporting license usage
linecount	Splunk Default field		
punct	Splunk Default field		
source	Splunk Default field		
sourcetype	Splunk Default field		

splunk_server	Splunk Default field		
splunk_server_group	Splunk Default field		
timeendpos	Common Extracted Fields		
timestamp	Indicates the timestamp	Timestamp	Reporting license usage
timestartpos	Common Extracted Fields		
utilization	Specifies the utilization	Integer	Reporting license usage
view	Specifies the DNS view	String	

## Summary Indexes

### Summary Indexes Frequency

The field frequencies of all fields for each summary index are as mentioned below:

Summary Index	Report	Frequency	Cron Schedule	Earliest Time	Latest Time
ib_dns_summary	si_dns_reclaimed_object_count_trend	At every 30 <sup>th</sup> minute from 21 through 59	21-59/30 * * * *	30m@m	60m@m
	si_dns_top_clients	At every 30 <sup>th</sup> minute from 2 through 59	2-59/30 * * * *	30m@m	60m@m
	si_dns_query_reply	At every 30 <sup>th</sup> minute from 18 through 59	18-59/30 * * * *	30m@m	60m@m
	si_top_servfail_received_queries	At every 30 <sup>th</sup> minute from 7 through 59	7-59/30 * * * *	30m@m	60m@m
	si_dns_response_latency_trend	At every 30 <sup>th</sup> minute from 20 through 59	20-59/30 * * * *	30m@m	60m@m
	si_dns_member_qps_trend_per_hour	At minute 34	34 * * * *	@h	-1h@h
	si_top_nxdomain_query	At every 30 <sup>th</sup> minute from 5 through 59	5-59/30 * * * *	30m@m	60m@m
	si_dns_member_qps_trend_per_day	Every day 32 minutes past midnight	32 0 * * *	@d	-1d@d
	si_dns_member_qps_trend	At every 30 <sup>th</sup> minute from 12 through 59	12-59/30 * * * *	30m@m	60m@m



	si_dns_requested_domain	At every 30 <sup>th</sup> minute from 4 through 59	4-59/30 * * * *	30m@m	60m@m
	si_dns_qps_trend	At every 30 <sup>th</sup> minute from 10 through 59	10-59/30 * * * *	30m@m	60m@m
	si_top_servfail_sent_queries	At every 30 <sup>th</sup> minute from 6 through 59	6-59/30 * * * *	30m@m	60m@m
	si_ddns_update	At every 30 <sup>th</sup> minute from 6 through 59	6-59/30 * * * *	30m@m	60m@m
	si_dns_cache_hit_ratio	At every 30 <sup>th</sup> minute from 8 through 59	8-59/30 * * * *	30m@m	60m@m
	si_top_timeout_queries	At every 30 <sup>th</sup> minute from 8 through 59	8-59/30 * * * *	30m@m	60m@m
	si_dns_rpz_hits	At every 10 <sup>th</sup> minute from 2 through 59	2-59/10 * * * *	10m@m	20m@m
	si_top_clients_per_domain	At every 30 <sup>th</sup> minute from 3 through 59	3-59/30 * * * *	30m@m	60m@m
ib_dhcp_summary	si_dhcp_message	At every 30 <sup>th</sup> minute from 14 through 59	14-59/30 * * * *	30m@m	60m@m
	si_dhcp_usage_trend	At 22 minutes past every 8 <sup>th</sup> hour	22 */8 * * * *	15m@m	495m@m
	si_dhcp_top_lease_client	At every 30 <sup>th</sup> minute from 16 through 59	16-59/30 * * * *	30m@m	60m@m
	si_devices_denied_an_ip_address	At every 30 <sup>th</sup> minute from 19 through 59	19-59/30 * * * *	30m@m	60m@m
	si_dhcp_range_utilization_trend	At 24 minutes past every 8 <sup>th</sup> hour	24 */8 * * * *	15m@m	495m@m
	si_dhcp_top_os_by_network	At every 30 <sup>th</sup> minute from 16 through 59	16-59/30 * * * *	30m@m	60m@m
ib_dtc_summary	si_dtc_response_distribution	At 37 minutes past every 6 <sup>th</sup> hour	37 */6 * * * *	10m@m	370m@m
	si_adns_resource_pool_availability	At 23 minutes past every 6 <sup>th</sup> hour	23 */6 * * * *	10m@m	370m@m
	si_smart_dns_resource_snmp	At 47 minutes past every 6 <sup>th</sup> hour	47 */6 * * * *	10m@m	370m@m

	si_smart_dns_resource_availability	At 47 minutes past every 6 <sup>th</sup> hour	47 */6 * * *	10m@m	370m@m
ib_system_summary	si_index_disk_usage	At 37 minutes past every 6 <sup>th</sup> hour	37 */6 * * *	10m@m	370m@m
	si_memory_utilization	At every 30 <sup>th</sup> minute from 26 through 59	26-59/30 * * *	30m@m	60m@m
	si_traffic_rate	At every 30 <sup>th</sup> minute from 28 through 59	28-59/30 * * *	30m@m	60m@m
	si_cpu_usage	At every 30 <sup>th</sup> minute	*/30 * * * *	30m@m	60m@m
ib_security_summary	si_dns_tunneling_activity	At every 3030 <sup>th</sup> th minute from 11 through 59	11-59/30 * * *	30m@m	60m@m

**Note:**

- **cron schedule** - cron time scheduled to execute a search
- **earliest time** - specifies the earliest time for a search
- **latest time** - specifies the latest time for a saved search

### Common fields in summary indexes

Splunk server adds the following fields to every event in each summary index.

Field Name	Description of the field	Values/Range	Remarks
info_max_time	The <b>info_*</b> fields are added to each event when you use the <b>addinfo</b> command. This command is primarily an internally-used component of Summary Indexing. Click <a href="#">here</a> for more information. The latest time boundary for the search.	Integer	Splunk added special field
info_min_time	Specifies the earliest time boundary for search	Integer	Splunk added special field
info_search_time	Specifies the time when search was initiated	Integer	Splunk added special field
search_name	Specifies the name of the saved search	Example: si-search-dns-query-reply	Splunk added special field
search_now	Specifies the time when search was scheduled to run	Integer	Splunk added special field

## Infoblox DNS Summary

Note: \*psrsvd\* stands for \*prestats reserved{\*}. Syntax is psrsvd\_[type]\_[fieldname]. These special fields are added by Splunk to summary index data that begins with \*psrsvd\* when you initiate search using the \*si\*\* command to populate a summary index. See [List of available psrsvd types](#) from Splunk docs.

Extracted Field Name	Description of the field	Reports	Values/Range	Source of Data	Remarks
CLIENT	Specifies the IP address of the DNS client		Example: 10.39.18.60		
COUNT	Specifies the count of DNS queries	si_dns_top_clients	Integer		
	Specifies the count of SERVFAIL errors that are received for DNS clients	si_top_servfail_received_queries	Integer		
	Specifies the count of NXDOMAIN/NOERROR replies for DNS clients	si_top_nxdomain_query	Integer		
	Specifies the count of DNS domain name requests	si_dns_requested_domain	Integer		
	Specifies the count of DNS queries per second	si_dns_qps_trend	Integer		
	Specifies the count of DNS SERVFAIL errors that are sent for DNS queries	si_top_servfail_sent_queries	Integer		
	Specifies the count of DNS timed-out recursive queries	si_top_timeout_queries	Integer		
	Specifies the average count of DNS RPX hits	si_dns_rpz_hits	Integer		
	Specifies the count of DNS clients per domain	si_top_clients_per_domain	Integer		
EA	Common Extracted Fields				
FQDN	Specifies the fully qualified domain name	si_dns_requested_domain and si_top_clients_per_domain	Example: 213.31.102.10.in-addr.arpa		
HWTYP	Common Extracted Fields				
MAX_DB_OBJECTS	Common Extracted Fields				

MAX_DHCP_LPS	Common Extracted Fields				
MAX_DNS_QPS	Common Extracted Fields				
MEMBER	Specifies the member		String	Infoblox DNS Summary	
MEMBER_IP	Common Extracted Fields				
TLD	Specifies top level domain names	si_dns_requested_domain	Example: arpa		
TYPE	Specifies the DNS response type	si_dns_query_reply, si_dns_qps_trend, and si_ddns_update	SUCCESS/ NOERROR OR REFERRAL OR NXRRSET OR NXDOMAIN OR REFUSED OR OTHER		
VIEW	It refers to the DNS view key to map DNS view through lookup. See display_name field.	si_dns_requested_domain, si_dns_top_clients, si_dns_member_qps_trend_per_hour, si_dns_member_qps_trend_per_day, si_dns_member_qps_trend, si_dns_qps_trend, si_ddns_update, si_dns_cache_hit_ratio, si_dns_rpz_hits, si_top_clients_per_domain, si_top_timeout_queries, si_top_servfail_sent_queries, si_top_nxdomain_query, and si_top_servfail_received_queries	Example: _default		
date_hour	Splunk Default field				
date_mday	Splunk Default field				
date_minute	Splunk Default field				
date_month	Splunk Default field				
date_second	Splunk Default field				
date_wday	Splunk Default field				
date_year	Splunk Default field				

date_zone	SplunkReporting Data Model				
display_name	Specifies the DNS view	si_dns_requested_domain, si_dns_top_clients, si_dns_member_qps_trend_per_hour, si_dns_member_qps_trend_per_day, si_dns_member_qps_trend, si_dns_qps_trend, si_ddns_update, si_dns_cache_hit_ratio, si_dns_rpz_hits, si_top_clients_per_domain, si_top_timeout_queries, si_top_servfail_sent_queries, si_top_nxdomain_query, and si_top_servfail_received_queries	Example: default.MS-2016		
eventtype	Splunk Default field				
host	Splunk Default field				
index	Splunk Default field				
info_max_time	Common summary index fields				
info_min_time	Common summary index fields				
info_search_time	Common summary index fields				
linecount	Splunk Default field				
orig_host	Specifies the host name of the data source		Example: infoblox.com		Splunk added default field
psrsvd_ct_COUNT	Here, ct = count. It contains the count information for the COUNT field.	si_dns_query_reply and si_dns_qps_trend			Splunk added special field
psrsvd_ct_LATENCY	Contains the count information for the LATENCY field	si_dns_response_latency_trend			Splunk added special field

psrsvd_ct_QCOUNT	Contains the count information for the QCOUNT field	si_dns_member_qps_tr end_per_hour, si_dns_member_qps_tr end_per_day, and si_dns_member_qps_tr end			Splunk added special field
psrsvd_gc	Here, gc = group count. It indicates the count for stats grouping and it is not scoped to a single field.	si_dns_query_reply, si_dns_response_latency_trend, si_dns_member_qps_tr end_per_hour, si_dns_member_qps_tr end_per_day, si_dns_member_qps_tr end, and si_dns_qps_trend			Splunk added special field
psrsvd_nc_COUNT	Here, nc = numerical count. It indicates the number of numerical values and contains the numerical count information for the COUNT field.	si_dns_query_reply and si_dns_qps_trend			Splunk added special field
psrsvd_nc_LATENCY	Contains the numerical count information for the LATENCY field	si_dns_response_latency_trend			Splunk added special field
psrsvd_nc_QCOUNT	Contains the numerical count information for the QCOUNT field	si_dns_member_qps_tr end_per_hour, si_dns_member_qps_tr end_per_day, and si_dns_member_qps_tr end			Splunk added special field
psrsvd_nx_QCOUNT	Here, nx = maximum numerical value. It contains the maximum numerical value information for the QCOUNT field.	si_dns_member_qps_tr end_per_hour and si_dns_member_qps_tr end_per_day			Splunk added special field
psrsvd_sm_COUNT	Here, sm = sum. It contains the sum information for the COUNT field.	si_dns_query_reply and si_dns_qps_trend			Splunk added special field
psrsvd_sm_LATENCY	Contains the sum information for the LATENCY field.	si_dns_response_latency_trend			Splunk added special field
psrsvd_sm_QCOUNT	Contains the sum information for the QCOUNT field	si_dns_member_qps_tr end_per_hour, si_dns_member_qps_tr end_per_day, and si_dns_member_qps_tr end			Splunk added special field

psrsvd_sx_QCOUNT	Here, sx = maximum lexicographical value. It contains the maximum lexicographical value information for the QCOUNT field	si_dns_member_qps_trend_per_hour and si_dns_member_qps_trend_per_day			Splunk added special field
psrsvd_v	Here, v = version. This is not scoped to a single field.	si_dns_query_reply, si_dns_response_latency_trend, si_dns_member_qps_trend_per_hour, si_dns_member_qps_trend_per_day, si_dns_member_qps_trend, and si_dns_qps_trend			Splunk added special field
psrsvd_vt_COUNT	Here, vt = value type. It contains precision of the associated field. This field contains precision of the COUNT field.	si_dns_query_reply and si_dns_qps_trend			Splunk added special field
psrsvd_vt_LATENCY	Contains precision of the LATENCY field	si_dns_response_latency_trend			Splunk added special field
psrsvd_vt_QCOUNT	Contains precision of the QCOUNT field	si_dns_member_qps_trend_per_hour, si_dns_member_qps_trend_per_day, and si_dns_member_qps_trend			Splunk added special field
<b>report</b>	<b>Contains the name of the report that populates the summary index</b>				
	<b>DNS Scavenge Object Count Trend data</b>	si_dns_reclaimed_object_count_trend			
	<b>DNS Top Clients report data</b>	si_dns_top_clients			
	<b>DNS Replies Trend data</b>	si_dns_query_reply			
	<b>DNS Top SERVFAIL Errors Received Report data</b>	si_top_servfail_received_queries			
	<b>DNS Response Latency Trend data</b>	si_dns_response_latency_trend			
	<b>DNS Daily Peak Hour Query Rate by Member Report data</b>	si_dns_member_qps_trend_per_hour			

	<b>DNS Top NXDOMAIN / NOERROR (no data) Report data</b>	si_top_nxdomain_query			
	<b>DNS Daily Query Rate by Member Report data</b>	si_dns_member_qps_trend_per_day			
	<b>DNS Query Rate by Member Report data</b>	si_dns_member_qps_trend			
	<b>DNS Top Requested Domain Names Report data</b>	si_dns_requested_domain			
	<b>DNS Queries Per Second Trend data</b>	si_dns_qps_trend			
	<b>DNS Top SERVFAIL Errors Sent Report data</b>	si_top_servfail_sent_queries			
	<b>DDNS Update Rate Trend data</b>	si_ddns_update			
	<b>DNS Cache Hit Rate Trend data</b>	si_dns_cache_hit_ratio			
	<b>DNS Top Timed-Out Recursive Queries Report data</b>	si_top_timeout_queries			
	<b>DNS RPZ Hits Reports data</b>	si_dns_rpz_hits			
	<b>DNS Top Clients per Domain Report data</b>	si_top_clients_per_domain			
search_name	Common summary index fields				
search_now	Common summary index fields				
source	Splunk Default field				
sourcetype	Splunk Default field				
splunk_server	Splunk Default field				
splunk_server_group	Splunk Default field				
timeendpos	Common Extracted Fields				



timestartpos	Common Extracted Fields				
--------------	-------------------------	--	--	--	--

## Infoblox DHCP Summary

Extracted Field Name	Description of the field	Reports	Values/Range	Source of Data	Remarks
ACTION	Specifies the action		String. Example: Issued	Infoblox DHCP summary	
DEVICE_CLASS	Specifies the device class		String. Example: Linux		
DHCP_RANGE	Specifies the DHCP range		Network range. Example: 10.0.0.1-10.0.0.20 0		
EA	Common Extracted fields				
FP	Specifies the fingerprint data		String. Example: No Match	Infoblox DHCP summary	
HWTYPE	Common Extracted Fields				
LEASED_IP	Specifies the lease IP address		IP address	Infoblox DHCP summary	
MAC_DUID	Specifies the MAC address		MAC address	Infoblox DHCP summary	
MAX_DB_OBJECTS	Common Extracted Fields				
MAX_DHCP_LPS	Common Extracted Fields				
MAX_DNS_QPS	Common Extracted Fields				
MEMBER_IP	Common Extracted Fields				
Protocol	Specifies the DHCP protocol		String. Example: IPV4	Infoblox DHCP summary	
SFP	Specifies the SFP		String. Example: Ubuntu/Debian 5/ Knoppix 6		

VIEW	It refers to the DNS view key to map the DNS view through lookup. See display_name field		String		
date_hour	<i>Splunk Default field</i>				
date_mday	Splunk Default field				
date_minute	Splunk Default field				
date_month	Splunk Default field				
date_second	Splunk Default field				
date_wday	Splunk Default field				
date_year	Splunk Default field				
date_zone	Splunk Default field				
dhcp_utilization_status	Specifies the DHCP utilization status		String	Infoblox DHCP summary	
display_name	Specifies the DNS view		String		
end_address	Specifies the end IP address		IP address	Infoblox DHCP summary	
eventtype	Splunk Default field				
host	Splunk Default field				
index	Splunk Default field				
info_max_time	Common summary index fields				

info_min_time	Common summary index fields				
info_search_time	Common summary index fields				
linecount	Splunk Default field				
members	Specifies the DHCP member		String. Example: infoblox.localdomain	Infoblox DHCP summary	
ms_servers	Specifies the MS servers		IP address	Infoblox DHCP summary	
orig_host	Specifies the host name of the data source		Example: infoblox.com		Splunk added default field
psrsvd_ct_FREE_ADDRESSES	Specifies the count information for FREE_ADDRESSES field	si_dhcp_usage_trend			Splunk added special field
psrsvd_ct_dhcp_utilization	Specifies the count for dhcp_utilization field	si_dhcp_range_utilization_trend			Splunk added special field
psrsvd_ct_dynamic_hosts	Specifies the count for dynamic_hosts field	si_dhcp_usage_trend			Splunk added special field
psrsvd_ct_static_hosts	Specifies the count for static_hosts field	si_dhcp_usage_trend			Splunk added special field
psrsvd_ct_v4ack	Specifies the count for v4ack field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v4decline	Specifies the count for v4decline field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v4discover	Specifies the count for v4discover field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v4inform	Specifies the count for v4inform field	si-search-dhcp-message			Splunk added special field

psrsvd_ct_v4leaseactive	Specifies the count for v4leaseactive field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v4leasequery	Specifies the count for v4leasequery field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v4leaseunassigned	Specifies the count for v4leaseunassigned field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v4leaseunknown	Specifies the count for v4leaseunknown field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v4nak	Specifies the count for v4nak field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v4offer	Specifies the count for v4offer field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v4release	Specifies the count for v4release field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v4request	Specifies the count for v4request field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6advertise	Specifies the count for v6advertise field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6confirm	Specifies the count for v6confirm field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6decline	Specifies the count for v6decline field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6information_request	Specifies the count for v6information_request field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6leasequery	Specifies the count for v6leasequery field	si-search-dhcp-message			Splunk added special field

psrsvd_ct_v6leasequery_reply	Specifies the count for v6leasequery_reply field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6rebind	Specifies the count for v6rebind field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6reconfigure	Specifies the count for v6reconfigure field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6relay_forward	Specifies the count for v6relay_forward field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6relay_reply	Specifies the count for v6relay_reply field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6release	Specifies the count for v6release field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6renew	Specifies the count for v6renew field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6reply	Specifies the count for v6reply field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6request	Specifies the count for v6request field	si-search-dhcp-message			Splunk added special field
psrsvd_ct_v6solicit	Specifies the count for v6solicit field	si-search-dhcp-message			Splunk added special field
psrsvd_gc	Here, gc = group count. The count for stats grouping and not scoped to a single field.				Splunk added special field
psrsvd_nc_FREE_ADDRESSES	Specifies the numerical count for FREE_ADDRESSES field	si_dhcp_usage_trend			Splunk added special field

psrsvd_nc_dhcp_utilization	Specifies the numerical count for dhcp_utilization field	si_dhcp_range_utilization_trend			Splunk added special field
psrsvd_nc_dynamic_hosts	Specifies the numerical count for dynamic_hosts field	si_dhcp_usage_trend			Splunk added special field
psrsvd_nc_static_hosts	Specifies the numerical count for static_hosts field	si_dhcp_usage_trend			Splunk added special field
psrsvd_nc_v4ack	Specifies the numerical count for v4ack field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v4decline	Specifies the numerical count for v4decline field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v4discover	Specifies the numerical count for v4discover field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v4inform	Specifies the numerical count for v4inform field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v4leaseactive	Specifies the numerical count for v4leaseactive field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v4leasequery	Specifies the numerical count for v4leasequery field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v4leaseunassigned	Specifies the numerical count for v4leaseunassigned field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v4leaseunknown	Specifies the numerical count for v4leaseunknown field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v4nak	Specifies the numerical count for v4nak field	si-search-dhcp-message			Splunk added special field

psrsvd_nc_v4offer	Specifies the numerical count for v4offer field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v4release	Specifies the numerical count for v4release field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v4request	Specifies the numerical count for v4request field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6advertise	Specifies the numerical count for v6advertise field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6confirm	Specifies the numerical count for v6confirm field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6decline	Specifies the numerical count for v6decline field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6information_request	Specifies the numerical count for v6information_request field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6leasequery	Specifies the numerical count for v6leasequery field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6leasequery_reply	Specifies the numerical count for v6leasequery_reply field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6rebind	Specifies the numerical count for v6rebind field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6reconfigure	Specifies the numerical count for v6reconfigure field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6relay_forward	Specifies the numerical count for v6relay_forward field	si-search-dhcp-message			Splunk added special field

psrsvd_nc_v6relay_reply	Specifies the numerical count for v6relay_reply field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6release	Specifies the numerical count for v6release field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6renew	Specifies the numerical count for v6renew field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6reply	Specifies the numerical count for v6reply field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6request	Specifies the numerical count for v6request field	si-search-dhcp-message			Splunk added special field
psrsvd_nc_v6solicit	Specifies the numerical count for v6solicit field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_FREE_ADDRESSES	Specifies the sum for FREE_ADDRESSES field	si_dhcp_usage_trend			Splunk added special field
psrsvd_sm_dhcp_utilization	Specifies the sum for dhcp_utilization field	si_dhcp_range_utilization_trend			Splunk added special field
psrsvd_sm_dynamic_hosts	Specifies the sum for dynamic_hosts field	si_dhcp_usage_trend			Splunk added special field
psrsvd_sm_static_hosts	Specifies the sum for static_hosts field	si_dhcp_usage_trend			Splunk added special field
psrsvd_sm_v4ack	Specifies the sum for v4ack field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v4decline	Specifies the sum for v4decline field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v4discover	Specifies the sum for v4discover field	si-search-dhcp-message			Splunk added special field



psrsvd_sm_v4inform	Specifies the sum for v4inform field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v4leaseactive	Specifies the sum for v4leaseactive field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v4leasequery	Specifies the sum for v4leasequery field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v4leaseunassigned	Specifies the sum for v4leaseunassigned field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v4leaseunknown	Specifies the sum for v4leaseunknown field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v4nak	Specifies the sum for v4nak field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v4offer	Specifies the sum for v4offer field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v4release	Specifies the sum for v4release field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v4request	Specifies the sum for v4request field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6advertise	Specifies the sum for v6advertise field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6confirm	Specifies the sum for v6confirm field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6decline	Specifies the sum for v6decline field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6information_request	Specifies the sum for v6information_request field	si-search-dhcp-message			Splunk added special field

psrsvd_sm_v6leasequery	Specifies the sum for v6leasequery field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6leasequery_reply	Specifies the sum for v6leasequery_reply field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6rebind	Specifies the sum for v6rebind field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6reconfigure	Specifies the sum for v6reconfigure field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6relay_forward	Specifies the sum for v6relay_forward field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6relay_reply	Specifies the sum for v6relay_reply field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6release	Specifies the sum for v6release field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6renew	Specifies the sum for v6renew field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6reply	Specifies the sum for v6reply field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6request	Specifies the sum for v6request field	si-search-dhcp-message			Splunk added special field
psrsvd_sm_v6solicit	Specifies the sum for v6solicit field	si-search-dhcp-message			Splunk added special field
psrsvd_v	Here, v = version. This is not scoped to a single field.	si_dhcp_usage_trend, si_dhcp_top_lease_client, si_dhcp_range_utilization_trend, si_dhcp_top_os_by_network, and si-search-dhcp-message			Splunk added special field

psrsvd_vt_FREE_ADDRESSES	Contains precision of the FREE_ADDRESSES field	si_dhcp_usage_trend			Splunk added special field
psrsvd_vt_dhcp_utilization	Contains precision of the dhcp_utilization field	si_dhcp_range_utilization_trend			Splunk added special field
psrsvd_vt_dynamic_hosts	Contains precision of the dynamic_hosts field	si_dhcp_usage_trend			Splunk added special field
psrsvd_vt_static_hosts	Contains precision of the static_hosts field	si_dhcp_usage_trend			Splunk added special field
psrsvd_vt_v4ack	Contains precision of the v4ack field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v4decline	Contains precision of the v4decline field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v4discover	Contains precision of the v4discover field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v4inform	Contains precision of the v4inform field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v4leaseactive	Contains precision of the v4leaseactive field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v4leasequery	Contains precision of the v4leasequery field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v4leaseunsigned	Contains precision of the v4leaseunassigned field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v4leaseunknown	Contains precision of the v4leaseunknown field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v4nak	Contains precision of the v4nak field	si-search-dhcp-message			Splunk added special field

psrsvd_vt_v4offer	Contains precision of the v4offer field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v4release	Contains precision of the v4release field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v4request	Contains precision of the v4request field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6advertise	Contains precision of the v6advertise field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6confirm	Contains precision of the v6confirm field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6decline	Contains precision of the v6decline field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6information_request	Contains precision of the v6information_request field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6leasequery	Contains precision of the v6leasequery field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6leasequery_reply	Contains precision of the v6leasequery_reply field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6rebind	Contains precision of the v6rebind field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6reconfigure	Contains precision of the v6reconfigure field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6relay_forward	Contains precision of the v6relay_forward field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6relay_reply	Contains precision of the v6relay_reply field	si-search-dhcp-message			Splunk added special field

psrsvd_vt_v6release	Contains precision of the v6release field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6renew	Contains precision of the v6renew field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6reply	Contains precision of the v6reply field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6request	Contains precision of the v6request field	si-search-dhcp-message			Splunk added special field
psrsvd_vt_v6solicit	Contains precision of the v6solicit field	si-search-dhcp-message			Splunk added special field
<b>report</b>	<b>Name of the report that is populating the summary index</b>				
	<b>DHCP Message Rate Trend data</b>	si-search-dhcp-message			
	<b>DHCPv4 Usage Trend data</b>	si_dhcp_usage_trend			
	<b>DHCP Top Lease Clients report data</b>	si_dhcp_top_lease_client			
	<b>Top Devices Denied an IP Address report data</b>	si_devices_denied_an_ip_address			
	<b>DHCPv4 Range Utilization Trend</b>	si_dhcp_range_utilization_trend			
	<b>Device and Device Classes reports data</b>	si_dhcp_top_os_by_network			
search_name	Common summary index fields				
search_now	Common summary index fields				
source	Splunk Default field				

sourcetype	Splunk Default field				
splunk_server	Splunk Default field				
splunk_server_group	Splunk Default field				
start_address	Specifies the start IP address		IP address	Infoblox DHCP summary	
timeendpos	Common Extracted Fields				
timestartpos	Common Extracted Fields				
View	Specifies the network view		String. Example: default	Infoblox DHCP summary	

### Infoblox DTC Summary

Extracted Field Name	Description of the field	Reports	Values/Range	Source of Data	Remarks
EA	Common Extracted Fields				
HWTYP	Common Extracted Fields				
MAX_DB_OBJECTS	Common Extracted Fields				
MAX_DHCP_LPS	Common Extracted Fields				
MAX_DNS_QPS	Common Extracted Fields				
MEMBER_IP	Common Extracted Fields				
date_hour	Splunk Default field				
date_mday	Splunk Default field				
date_minute	Splunk Default field				

date_month	Splunk Default field				
date_second	Splunk Default field				
date_wday	Splunk Default field				
date_year	Splunk Default field				
date_zone	Splunk Default field				
eventtype	Splunk Default field				
host	Splunk Default field				
index	Splunk Default field				
info_max_time	Common summary index fields				
info_min_time	Common summary index fields				
info_search_time	Common summary index fields				
linecount	Splunk Default field				
Monitor	Specifies the monitor		String. Example: https	Infoblox DTC summary	
orig_host	Specifies the host name of the data source		Example: infoblox.com		Splunk added default field
pool	Specifies the Pool		String. Example: Pool	Infoblox DTC summary	
psrsvd_ct_availability	Specifies the count information for available field	si_adns_resource_pool_availability and si_smart_dns_resource_availability			Splunk added special field

psrsvd_ct_respons e_count	Specifies the count information for response_count field	si_dtc_response_distri bution			Splunk added special field
psrsvd_ct_unavaila ble	Specifies the count information for unavailable field	si_adns_resource_poo l_availability and si_smart_dns_resourc e_availability			Splunk added special field
psrscd_ct_value	Specifies the count information for value field	si_smart_dns_resourc e_snmp			Splunk added special field
psrsvd_gc	Here, gc = group count. This is the count for stats grouping and it is not scoped to a single field.	si_dtc_response_distri bution, si_smart_dns_resourc e_snmp, si_adns_resource_poo l_availability, and si_smart_dns_resourc e_availability			Splunk added special field
psrsvd_nc_availabl e	Specifies the numerical count information for available field	si_adns_resource_poo l_availability and si_smart_dns_resourc e_availability			Splunk added special field
psrsvd_nc_respons e_count	Specifies the numerical count information for response_count field	si_dtc_response_distri bution			Splunk added special field
psrsvd_nc_unavail able	Specifies the numerical count information for unavailable field	si_adns_resource_poo l_availability and si_smart_dns_resourc e_availability			Splunk added special field
psrsvd_nc_value	Specifies the numerical count information for value field	si_smart_dns_resourc e_snmp			Splunk added special field
psrsvd_sm_availab le	Specifies the sum information for available field	si_adns_resource_poo l_availability and si_smart_dns_resourc e_availability			Splunk added special field
psrsvd_sm_respon se_count	Specifies the sum information for response_count field	si_dtc_response_distri bution			Splunk added special field



psrsvd_sm_unavailable	Specifies the sum information for unavailable field	si_adns_resource_pool_availability and si_smart_dns_resource_availability			Splunk added special field
psrsvd_sm_value	Specifies the sum information for value field	si_smart_dns_resource_snmp			Splunk added special field
psrsvd_v	Here, v = version. This is not scoped to a single field.	si_dtc_response_distribution, si_smart_dns_resource_snmp, si_adns_resource_pool_availability, and si_smart_dns_resource_availability			Splunk added special field
psrsvd_vt_available	Contains precision of the available field	si_adns_resource_pool_availability and si_smart_dns_resource_availability			Splunk added special field
psrsvd_vt_response_count	Contains precision of the response_count field	si_dtc_response_distribution			Splunk added special field
psrsvd_vt_unavailable	Contains precision of the unavailable field	si_adns_resource_pool_availability and si_smart_dns_resource_availability			Splunk added special field
psrsvd_vt_value	Contains precision of the value field	si_smart_dns_resource_snmp			Splunk added special field
<b>report</b>	<b>Name of the report that populates the summary index</b>				
	<b>DNS Traffic Control Response Distribution Trend data</b>	si_dtc_response_distribution			
	<b>DNS Traffic Control Resource Pool Availability reports data</b>	si_adns_resource_pool_availability			
	<b>DNS Traffic Control Resource SNMP reports data</b>	si_smart_dns_resource_snmp			

	<b>DNS Traffic Control Resource Availability reports data</b>	si_smart_dns_resource_availability			
resource	Specifies the resource		String. Example: Server	Infoblox DTC summary	
search_name	Common summary index fields				
search_now	Common summary index fields				
source	Splunk Default field				
sourcetype	Splunk Default field				
splunk_server	Splunk Default field				
splunk_server_group	Splunk Default field				
timeendpos	Common Extracted Fields				
timestartpos	Common Extracted Fields				

### Infoblox System Summary

Extracted Field Name	Description of the field	Reports	Values/Range	Source of Data	Remarks
EA	Common Extracted Fields				
HWTYP	Common Extracted Fields				
MAX_DB_OBJECTS	Common Extracted Fields				
MAX_DHCP_LPS	Common Extracted Fields				
MAX_DNS_QUEUES	Common Extracted Fields				

MEMBER	Specifies the member		String. Example: infoblox.localdomain : inbound		
MEMBER_IP	Common Extracted Fields				
date_hour	Splunk Default field				
date_mday	Splunk Default field				
date_minute	Splunk Default field				
date_month	Splunk Default field				
date_second	Splunk Default field				
date_wday	Splunk Default field				
date_year	Splunk Default field				
date_zone	Splunk Default field				
eventtype	Splunk Default field				
host	Splunk Default field				
index	Splunk Default field				
info_max_time	Common summary index fields				
info_min_time	Common summary index fields				
info_search_time	Common summary index fields				
linecount	Splunk Default field				
orig_host	Specifies the host name of the data source		Example: infoblox.com		Splunk added default field
psrsvd_ct_CPU_PERCENT	Specifies the count information for the CPU_PERCENT field	si_cpu_usage			Splunk added special field
psrsvd_ct_MEMORY_PERCENT	Specifies the count information for the MEMORY_PERCENT field	si_memory_utilization			Splunk added special field

psrsvd_ct_TRA F_VALUE	Specifies the count information for TRAF_VALUE field	si_traffic_rate			Splunk added special field
psrsvd_gc	Here, gc = group count. This is the count for a stats grouping and it is not scoped to a single field.	si_memory_utilization, si_traffic_rate, and si_cpu_usage			Splunk added special field
psrsvd_nc_CPU _PERCENT	Specifies the numerical count information for CPU_PERCENT field	si_cpu_usage			Splunk added special field
psrsvd_nc_ME MORY_PERCE NT	Specifies the numerical count information for MEMORY_PERCENT field	si_memory_utilization			Splunk added special field
psrsvd_nc_TRA F_VALUE	Specifies the numerical count information for TRAF_VALUE field	si_traffic_rate			Splunk added special field
psrsvd_sm_CP U_PERCENT	Specifies the sum for CPU_PERCENT field	si_cpu_usage			Splunk added special field
psrsvd_sm_ME MORY_PERCE NT	Specifies the sum for MEMORY_PERCENT field	si_memory_utilization			Splunk added special field
psrsvd_sm_TRA F_VALUE	Specifies the sum for TRAF_VALUE field	si_traffic_rate			Splunk added special field
psrsvd_v	Here, v = version. This is not scoped to a single field.	si_memory_utilization, si_traffic_rate, and si_cpu_usage			Splunk added special field
psrsvd_vt_CPU _PERCENT	Contains precision of the CPU_PERCENT field	si_cpu_usage			Splunk added special field
psrsvd_vt_MEM ORY_PERCEN T	Contains precision of the MEMORY_PERCENT field	si_memory_utilization			Splunk added special field
psrsvd_vt_TRA F_VALUE	Contains precision of the TRAF_VALUE field	si_traffic_rate			Splunk added special field
<b>report</b>	<b>Specifies the name of the report that is populating the summary index</b>				
	<b>Index Disk Usage Report Data</b>	si_index_disk_usage			
	<b>Memory Utilization Trend data</b>	si_memory_utilization			

	<b>Traffic Rate by Member report data</b>	si_traffic_rate			
	<b>CPU Utilization Trend data</b>	si_cpu_usage			
search_name	Common summary index fields				
search_now	Common summary index fields				
source	Splunk Default field				
sourcetype	Splunk Default field				
splunk_server	Splunk Default field				
splunk_server_group	Splunk Default field				
timeendpos	Common Extracted Fields				
timestartpos	Common Extracted Fields				

### Infoblox Security Summary

Extracted Field Name	Description of the field	Reports	Values/Range	Source of Data	Remarks
ACTIVE_COUNT	Specifies the active count		Integer	Infoblox security summary	
BLOCK_END	Specifies the block end IP address		Integer	Infoblox security summary	
BLOCK_START	Specifies the block start IP address		Integer	Infoblox security summary	
DNST_CATEGORY	Specifies the destination category		String		
EA	Common Extracted Fields				
HWTYP	Common Extracted Fields				
MAX_DB_OBJECTS	Common Extracted Fields				
MAX_DHCP_LEASES	Common Extracted Fields				

MAX_DNS_QPS	Common Extracted Fields				
MEMBER_IP	Common Extracted Fields				
NAT_STATUS	Specifies the NAT status		String	Infoblox security summary	
RULE_DESCRIPTION	Specifies the rule description		String. Example: This rule drops unexpected OSPF packets when OSPF is disabled.		
RULE_NAME	Specifies the rule name		String. Example: DROP OSPF unexpected		
RULE_SID	Specifies the rule SID		Integer	Infoblox security summary	
SOURCE_IP	Specifies the source IP		IP address	Infoblox security summary	
SOURCE_PORT	Specifies the source port		Integer	Infoblox security summary	
date_hour	Splunk Default field				
date_mday	<i>Splunk Default field</i>				
date_minute	Splunk Default field				
date_month	Splunk Default field				
date_second	Splunk Default field				
date_wday	Splunk Default field				
date_year	Splunk Default field				
date_zone	Splunk Default field				
eventtype	Splunk Default field				
host	Splunk Default field				
index	Splunk Default field				

info_max_time	Common summary index fields				
info_min_time	Common summary index fields				
info_search_time	Common summary index fields				
linecount	Splunk Default field				
orig_host	Specifies the host name of the data source		Example: infoblox.com		Splunk added default field
<b>report</b>	<b>Name of the report that is populating the summary index</b>				
	<b>DNS Tunneling Activity Reports data</b>	si_dns_tunneling_activity			
search_name	Common summary index fields				
search_now	Common summary index fields				
source	Splunk Default field				
sourcetype	Splunk Default field				
splunk_server	Splunk Default field				
splunk_server_group	Splunk Default field				
timeendpos	Common Extracted Fields				
timestartpos	Common Extracted Fields				

## Managing Reporting Data

You can do the following to manage the Reporting and Analytics App:

- [Upgrading Reporting and Analytics App](#)
- [Backing Up and Restoring the Infoblox Reporting and Analytics App](#)
- [Best Practices for Backing Up Reporting Data](#)
- [Reports with Data Synchronized from Microsoft Servers](#)

## Upgrading Reporting and Analytics App



### Note

All the updates for Infoblox Reporting & Analytics App are available with the NIOS upgrade. Therefore, as a part of NIOS upgrade, no further steps are required for upgrading the Infoblox Reporting & Analytics App.

You can download the latest version of the Infoblox Reporting and Analytics App from the Infoblox Support site. You can then upgrade this App on your reporting server.

To upgrade the Reporting and Analytics App:

1. Check your current version of the App in the **Reporting Help** tab.
2. Check if there is a later version that is available on the Infoblox Support site.
3. Download the `.bin2` file from Infoblox Support site.
4. From the **Administration** tab, select the **Reporting** tab.
5. Click **Upgrade Reporting & Analytics App** from the Toolbar.
6. In the *Upgrade Reporting & Analytics App* dialog box, click **Select**. In the *Upload* dialog box, click **Select**, navigate to the `.bin2` file, select it, and then click **Upload**.
7. After the Reporting and Analytics App is upgraded, click **Restart** to restart the Reporting service.

## Backing Up and Restoring the Infoblox Reporting and Analytics App

To back up the Reporting and Analytics App, go to the **Grid** tab, select **Backup** -> **Grid Backup** -> **Manual Backup** from the Toolbar, and then select **Infoblox Reporting & Analytics App** to back up the app. When you back up the Reporting and Analytics App, the backup file is a `.bak` file that contains the reporting settings configured in the Grid Reporting Properties.

To restore the Reporting and Analytics App, go to the **Grid** tab, select **Restore** -> **Restore Grid** from the Toolbar, and then select **Infoblox Reporting & Analytics App** to restore the app.

### Backing Up Reporting Data

Before you back up the reporting database, ensure that the reporting service is enabled on the reporting server. You cannot perform or schedule a backup if the reporting service is disabled on the reporting server. If you want to upgrade your reporting server, back up all the data before you power down the server. During an upgrade, the reporting server is automatically upgraded after the Grid Master. You cannot control or schedule when to upgrade the reporting server. For information about upgrades and upgrade groups, see [Managing Upgrade Groups](#).

Note that reporting data backups are incremental backups, which means that backup files are copied to the designated file server only when there are new events generated since the last backup. Backing up of the reporting database to an FTP or SCP server using IPv4 or IPv6 is supported. The backup file is a `.tar.gz` file that contains the reporting data.

If you stop an ongoing backup process, backup files are still being copied to the designated file server. When you perform a subsequent backup, the appliance appends incremental data to these backup files.

You can manually back up the reporting database or schedule a backup, but you cannot perform both at the same time. The backup process starts when the indexed data rolls from the hot bucket to the warm bucket. The hot bucket includes all inbound events and actively written data. Indexed data moves to the warm bucket when one of the following conditions is met:

- The size of the reporting data reaches 1 GB
- Data is 90 days old
- The reporting server restarts

You can perform the following reporting data backups:

- Manual backups, as described in [Backing Up the Reporting Database Manually](#) below.
- Scheduled backups, as described in [Scheduling the Backup of the Reporting Database](#) below.



## Backing Up the Reporting Database Manually

- From the **Grid** tab, select **Backup** -> **ReportingBackup** -> **ManualBackup** from the Toolbar.
- In the *ManualReportingBackup* editor, complete the following:
  - **Status**: Displays the status of the backup process of the last operation.
  - **Backupto**: Select the destination of the backup file from the drop-down list:
    - **FTP**: Back up the reporting database to an FTP server.
      - **Filepath**: Enter the directory path. For example, you can enter `/archive/backups/Infoblox/`.
      - **IP Address of FTP Server**: The IP address of the FTP server.
      - **Username**: Enter the username of your FTP account.
      - **Password**: Enter the password of your FTP account.
    - **SCP**: Back up the reporting database to an SSH server that supports SCP.
      - **Filepath**: Enter the directory path. For example, you can enter `/archive/backups/Infoblox/`.
      - **IP Address of SCP Server**: The IP address of the SCP server.
      - **Username**: Enter the username of your SCP account.
      - **Use Keys**: If you select this checkbox, you can back up files to SCP without entering the password. The first time you select the checkbox, you need to enter the password. However, during subsequent times, the Infoblox server verifies whether Infoblox keys are available on the SCP server. If they are available, you can click the **Backup** button without entering the password. If Infoblox keys are not available on the SCP server, the following message is displayed:

```
Reporting backup has failed.
```
      - **Password**: Enter the password of your SCP account.
      - **Keys Type**: Select the SSH key type to be uploaded. At present, only ECDSA and RSA keys are supported. Click **Upload Keys** to upload the keys to the SCP server. If the keys are not available, click **Download Keys** to download the keys and manually add them to the SCP server.

### Notes

- If you are using Fedora, ECDSA keys are supported only on Fedora versions later than Fedora 12.
- When you select **FTP** or **SCP**, ensure that you have a valid user name and password on the server prior to backing up the files. Also ensure that the target SSH server has the required permissions for an SCP backup. The permission must be 755 and the target server must have write permission to the directory to which you upload the backup file.
- For an SCP backup, ensure that you are logged in as the user for whom the key was created. Also ensure that the `.ssh` directory on the server and the files it contains, have the correct permissions: 

```
chmod 600 ~/.ssh/authorized_keys && chmod 700 ~/.ssh/
```
- If you promote a Grid Master or perform an HA failover, you must upload the SSH key once again for a successful SCP backup using keys.

## Scheduling the Backup of the Reporting Database

- From the **Grid** tab, select **Backup** -> **ReportingBackup** -> **ScheduleBackup** from the Toolbar.
- In the *ScheduleReportingBackup* editor, complete the following:
  - **Status**: Displays the status of the backup process of the last operation. Select the destination of the backup file from the **Backupto** drop-down list:
  - **FTP**: Back up the reporting database files to an FTP server.
    - **IP Address of FTPServer**: The IP address of the FTP server.
    - **Directory Path**: Enter the directory path. For example, you can enter `/archive/backups`. The directory path cannot contain spaces. The folder or directory you enter here must already exist on the specified server. Do not include the file name in the directory path.
    - **Username**: Enter the username of your FTP account.
    - **Password**: Enter the password of your FTP account.

- **Recurrence:** Select how often you want to back up the files. You can select **Weekly**, **Daily**, or **Hourly** from the drop-down list. When you select **Weekly**, complete the following:
  - **Every:** Choose a day of the week from the drop-down list.
  - **Time:** Enter a time in the hh:mm:ss AM/PM format. You can also click the clock icon and select a time from the drop-down list. The Grid Master creates a backup file on the selected day and time every week.  
When you select **Daily**, enter a time in the hh:mm:ss AM/PM format. You can also select a time from the drop-down list.  
When you select **Hourly**, complete the following:
    - **Minutes after the Hour:** Enter the minute after the hour when the Grid Master creates a backup file. For example, enter 5 if you want the Grid Master to create a backup file five minutes after the hour every hour.
- **Disable Scheduled Backup:** Select this if you want to disable automatic backups from occurring now, but want to save the settings for future use.
- **SCP:** Back up the reporting database to an SSH server that supports SCP.
  - **IP Address of SCP Server:** The IP address of the SCP server.
  - **Directory Path:** Enter the directory path of the file. For example, you can enter `/archive/backups`. The directory path cannot contain spaces. The folder or directory you enter here must already exist on the specified server. Do not include the file name in the directory path.
  - **Username:** Enter the username of your SCP account.
  - **Use Keys:** If you select this checkbox, you can back up files to SCP without entering the password. The first time you select the checkbox, you need to enter the password. However, during subsequent times, the Infoblox server verifies whether Infoblox keys are available on the SCP server. If they are available, you can click the **Backup** button without entering the password. If Infoblox keys are not available on the SCP server, the following message is displayed:  
`Reporting backup has failed.`
  - **Password:** Enter the password of your SCP account.
  - **Keys Type:** Select the SSH key type to be uploaded. At present, only ECDSA and RSA keys are supported. Click **Upload Keys** to upload the keys to the SCP server. If the keys are not available, click **Download Keys** to download the keys and manually add them to the SCP server.
  - **Recurrence:** Select how often the scheduled backups should occur. You can select **Weekly**, **Daily**, or **Hourly**. For information, see the FTP section.
  - **Disable Scheduled Backup:** Select this if you want to disable automatic backups from occurring now. You can still save the settings for future use.

#### Notes

- If you are using Fedora, ECDSA keys are supported only on Fedora versions later than Fedora 12.
- When you select **FTP** or **SCP**, ensure that you have a valid user name and password on the server prior to backing up the files. Also ensure that the target SSH server has the required permissions for an SCP backup. The permission must be 755 and the target server must have write permission to the directory to which you upload the backup file.
- For an SCP backup, ensure that you are logged in as the user for whom the key was created. Also ensure that the `.ssh` directory on the server and the files it contains, have the correct permissions: `chmod 600 ~/.ssh/authorized_keys && chmod 700 ~/.ssh/`
- If you promote a Grid Master or perform an HA failover, you must upload the SSH key once again for a successful SCP backup using keys.

## Restoring the Reporting Database

Restoring the reporting database may take a long time to perform, and the reporting service is unavailable during a restore. Ensure that you back up the reporting database before you perform the operation. Restoring of reporting database from an FTP or SCP server using IPv4 or IPv6 is supported.

Note the following during a restore:

- The reporting service is unavailable.
- Existing reporting data is removed from the reporting server.

- Backup data is restored up to the amount the reporting server can accommodate. The Volume Used Today displayed in the Device Information section will not be updated after restoring the data. Also, when you restore data or execute the CLI command `reset reporting_data`, the volume violation count will be reset to zero on the second day.
1. From the **Grid** tab, select **Restore** -> **Restore Reporting** from the Toolbar.
  2. In the *Restore* dialog box, complete the following:
    - **Status**: Displays the status of the restore process, if in progress. Select the destination of the backup file from the **Restore from** drop-down list:
    - **FTP**: Restore the reporting backup files from an FTP server.
      - **Filepath**: Enter the directory path. For example, you can enter `/archive/backups/Infoblox/`.
      - **IP Address of FTP Server**: The IP address of the FTP server.
      - **Username**: Enter the username of your FTP server account.
      - **Password**: Enter the password of your FTP server account.
    - **SCP**: Restore the reporting backup files from a SCP server.
      - **Filepath**: Enter the directory path. For example, you can enter `/archive/backups/Infoblox/`.
      - **IP Address of SCP Server**: The IP address of the SCP server.
      - **Username**: Enter the username of your SCP server account.
      - **Password**: Enter the password of your SCP server account.
  3. Click **Restore**.

## Best Practices for Backing Up Reporting Data

The following are some best practices that you must implement when backing up reporting data:

### Best Practices for Backing Up

- The first backup of the reporting data is a full backup. Subsequent backups irrespective of scheduled or manual are incremental backups. Infoblox recommends that you do not change the backup directory for each backup.
- Reporting data backups are incremental. Any modification or deletion of the backed up data can create an impact when the data is restored.
- If you upgrade or replace the Reporting server, Infoblox recommends that you back up the reporting data and the Infoblox Reporting and Analytics App.
- Once the reporting data is restored on a Grid, subsequent backups include only the new data.

### Best Practices for Upgrading or Replacing the Reporting Server (Single Indexer)

- After you upgrade or replace the Reporting server, restore the Infoblox Reporting and Analytics App to restore any custom reports or Grid settings.
- After you upgrade or replace the Reporting server, restore the reporting data from the reporting backup to view the historic reporting data.

### Best Practices for Upgrading or Replacing the Reporting Server (Single or Multi-Site Cluster)

- If the Grid has a reporting configuration updated from a single indexer to a single or multi-site cluster, data replication starts only for the data indexed after clustering. If you upgrade or replace the old reporting server, you may lose the data indexed before clustering. Infoblox recommends that you restore the reporting backup to view historic data.

## Reports with Data Synchronized from Microsoft Servers

**Note:** The DNS reports listed in the following table displays data synchronized from the Microsoft servers only when you have enabled synchronization of reporting data for the Grid or the Microsoft servers. For information about enabling synchronization of DNS reporting data from the Microsoft server, see [Synchronizing DNS Reporting Data](#).

Infoblox supports the following versions of Microsoft Windows servers in displaying reporting data from both NIOS and the Microsoft servers:

- DNS Reports: Microsoft Windows 2012 R2, Microsoft Windows 2016 and Microsoft Windows 2019.
- DHCP Reports: Microsoft Windows 2008, Microsoft Windows 2008 R2, Microsoft Windows 2012, Microsoft Windows 2012 R2, Microsoft Windows 2016 and Microsoft Windows 2019.
- IPAM Reports: Microsoft Windows 2008, Microsoft Windows 2008 R2, Microsoft Windows 2012, Microsoft Windows 2012 R2, Microsoft Windows 2016 and Microsoft Windows 2019.

The following reports display data from both NIOS and the Microsoft servers. For detailed information about these reports, see [Predefined Dashboards](#).

### DNS Reports

- *DNS Top Requested Domain Names*
- *DNS Top Clients*
- *DNS Top Clients Per Domain*
- *DNS Query Rate by Query Type*
- *DNS Query Rate by Member*
- *DNS Daily Query Rate by Member*
- *DNS Daily Peak Hour Query Rate by Member*
- *DNS Top NXDOMAIN /NOERROR (no data)*
- *DNS Top SERVFAIL Errors Sent*
- *DNS Top SERVFAIL Errors Received*
- *DNS Top Timed-out Recursive Queries*
- *DNS Query Trend per IP Block Group*
- *DDNS Update Rate Trend*

Following DNS reports are generated if a Data Collector VM is registered with the Grid:

- *DNS Domain Queried by Client*
- *DNS Domain Query Trend*
- *DNS Top Clients by Query Type*
- *DNS Top Clients Querying MX Records*

### DHCP Reports

- *DHCPv4 Usage Trend*
- *DHCPv4 Usage Statistics*
- *DHCPv4 Range Utilization Trend*
- *DHCPv4 Top Utilized Networks*
- *DHCP Lease History*
- *DHCP Top Lease Clients*
- *DHCP Message Rate Trend*

### IPAM Reports

- *IPAMv4 Network Usage Statistics*
- *IPAMv4 Network Usage Trend*
- *IPAMv4 Top Utilized Networks*

## Infoblox Infrastructure Security

The mission-critical DNS infrastructure can become a vulnerable component in your network when it is inadequately protected by traditional security solutions and consequently used as an attack surface. Compromised DNS services can result in catastrophic network and system failures. To fully protect your network in today's cyber security threat environment, Infoblox sets a new DNS security standard by offering scalable, enterprise-grade, and integrated protection for your DNS infrastructure.

While your external (internet-facing) DNS server can be subject to cyber attacks such as DNS DDoS (Distributed Denial of Service) and others, threats can also come from the inside of your firewalls. Today's targeted attacks pose risk to both data and infrastructure inside an enterprise. You could have an endpoint infected with malware or threats trying to communicate with C&C (Command-and-Control) servers that use DNS as a protocol. You could also have a malicious

insider trying to steal valuable digital assets by opening a DNS tunnel or embedding data in DNS queries. Depending on how you want to protect your mission-critical DNS infrastructure, you can configure your Infoblox appliance to mitigate against external, internal, or both (external and internal) DNS threats.

This section contains information about the Infoblox infrastructure security features that protect external DNS from cyber DNS attacks and internal DNS from infrastructure attacks, data exfiltration, and APTs (Advanced Persistent Threats) and malware. It covers the following topics:

- [About Infoblox Advanced DNS Protection](#)
- [Infoblox DNS Firewall](#)
- [Infoblox Threat Insight](#)
- [Cisco ISE Integration](#)
- [Ecosystem - Outbound Notifications](#)

## Infoblox Advanced DNS Protection

The Infoblox Advanced DNS Protection solution employs hardware-accelerated security rules to detect, report upon, and stop attacks such as DDoS, DNS reflection, DNS amplification, DNS hijacking, and other network attacks targeting DNS authoritative applications. This security solution helps minimize "false positives" and ensures that your mission-critical DNS services continue to function even when under attack. For more information, see [Infoblox Advanced DNS Protection](#).

## Infoblox DNS Firewall

Infoblox DNS Firewall uses DNS RPZs (Response Policy Zones) for allowing reputable sources to dynamically communicate reputation domain names so you can implement policy controls for DNS lookups. For more information, see [Infoblox DNS Firewall](#).

## Infoblox Threat Insight

The Infoblox Threat Insight solution defends against data exfiltration through DNS tunneling for ultimate network protection. For more information, see [Infoblox Threat Insight](#).

## Security Ecosystem

The Infoblox security ecosystem comprises FireEye integrated RPZs for detecting malware and APTs and the TAXII (Trusted Automated eXchange of Indicator Information) service for mitigating cyber attacks. For more information, see [Infoblox DNS Firewall](#).

For best practices in securing your networks, you can also set up DNS blacklists or configure a security banner. When you enable DNS Integrity Check for top-level authoritative zones, the appliance verifies DNS data in the NS RRsets and glue records, and reports any data discrepancies so you can mitigate possible DNS domain hijacking. Following are other DNS security features for your network security:

## Access Control (Named ACLs)

To effectively manage your core network services, you can grant legitimate hosts access to specific operations on the appliance using an ACL (access control list) or anonymous ACEs (access control entries). You can also configure a named ACL and apply it to multiple operations, such as file distribution and DNS zone transfers. For more information, see [Configuring Access Control](#).

## DNS blacklists

Your organization can prevent customers or employees from accessing certain Internet resources, particularly web sites, by prohibiting a recursive DNS member from resolving queries for domain names that you specify. You can configure a recursive DNS member to redirect the DNS client to predefined IP addresses or return a REFUSED response code (indicating that resolution is not performed because of local policy), depending on the domain name. For more information, see [Blacklists](#).

## Security Banner

You can configure and publish a notice and consent banner as the first login screen that includes specific terms and conditions you want end users to accept before they log in to the Infoblox Grid. When you enable the notice and consent banner, users must accept the terms and conditions displayed on the consent screen before accessing the login screen of Grid Manager. For more information about configuring notice and consent banner, see [Managing a Grid](#).

## DNS Integrity Check

DNS domain hijacking or domain theft is the act of changing the registration of a domain name without the permission of its original registrant. In some cases, hijackers change the DNS data of a domain after gaining control of it. They consequently redirect users to a fraudulent site, instead of the legitimate site, on the Internet. To protect your authoritative DNS servers against this type of domain hijacking, you can configure the appliance to monitor NS records and glue records for top-level authoritative zones. Based on your configuration, the appliance periodically checks the DNS data for the zones and compares the data with that in the appliance database. The severity in data discrepancies can help identify possible domain hijacking. For more information about this feature, see [Configuring DNS Integrity Check for Authoritative Zones](#).

## About Infoblox Advanced DNS Protection

The Infoblox Advanced DNS Protection solution employs threat protection rules to detect, report upon, and stop DoS (Denial of Service), DDoS (Distributed Denial of Service) and other network attacks targeting DNS authoritative applications. Infoblox Advanced DNS Protection helps minimize "false positives" and ensures that your mission-critical DNS services continue to function even when under attack. For information about possible DNS threats, see [DNS and Network-Flood Threats](#).

You can deploy the Advanced DNS Protection solution on hardware-accelerated appliances (physical appliances only) as well as software-based appliances (both physical and virtual) in the Grid. Depending on the appliances you deploy, you must install applicable hardware-based licenses or Software ADP subscription licenses. For information about supported Infoblox appliances for Advanced DNS Protection and the applicable licenses, see [Supported Threat Protection Appliances and Licensing Requirements](#).

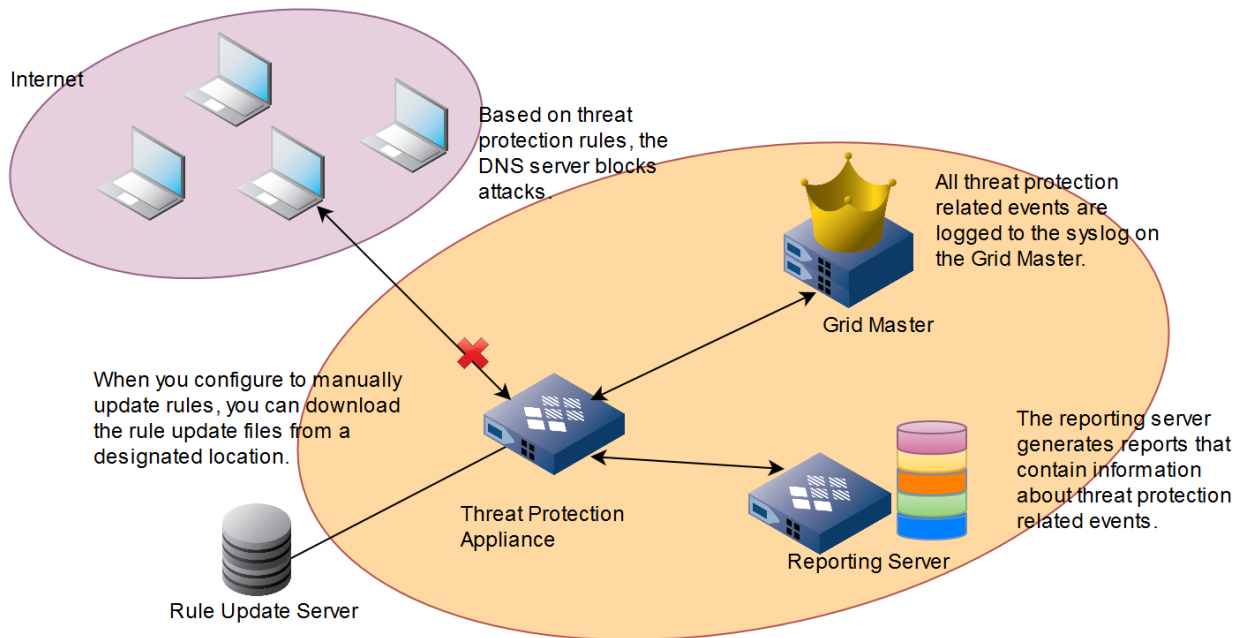
Infoblox Advanced DNS Protection is designed to provide visibility and protection against network floods and DNS attacks. It detects DNS attacks through predefined and custom threat protection rules, and mitigates DNS threats by dropping problematic packets while responding only to legitimate traffic. With valid licenses installed, you can subscribe to automatic rule updates that deliver near real-time protection against new and emerging attacks. You may also manually perform the rule update process based on your configuration. For information about threat protection rules, see [Understanding Threat Protection Rulesets and Rules](#).

Infoblox Advanced DNS Protection supports a set of predefined threat protection rules that detect and mitigate possible DNS threats. You can modify some of the parameters and assign actions such as logging events and applying mitigation to these rules. You can also create custom rules to suit your security needs. For more information, see [Understanding Threat Protection Rulesets and Rules](#).

As illustrated in the figure below, the threat protection appliance, acting as an authoritative DNS server, is added to the Grid. After installing valid threat protection licenses and configuring the appliance to serve as an Advance Appliance, it can now detect DNS threats and mitigate DNS threats based on threat protection rules. All threat protection related events, conformed to CEF (Common Event Format), are logged in the syslog on the Grid Master. To perform further

investigation about possible threats, the reporting server generates specific threat protection related reports. For information about how to monitor threat protection related events and reports, see [Monitoring Threat Protection Events](#).

### Infoblox Advanced DNS Protection Solution



## Limitations for Threat Protection Appliances

Hardware-based appliances support all existing DNS features (including HA support) that are applicable to DNS caching and authoritative applications, except the following:

- Configuration of multiple interfaces on the same subnet
- 10/100-Mbps gigabit Ethernet mode and fixed speed/duplex settings



### Note

The MGMT interface must be configured and Infoblox strongly recommends that the MGMT interface be used for Grid traffic.

Consider the following when the threat protection service is enabled on advanced appliances:

### For Hardware ADP

- Protected interfaces (LAN1 and LAN2) are limited to DNS traffic, protocols in support of DNS anycast (BGP and OSPF) and the standard IP protocols such as ICMP, as well as connections to NTP servers.
- The MGMT interface is used for other traffic, such as Grid, SSH, SNMP, NTP, and it will not be protected against DDoS attacks.
- You cannot run other services, such as FTP, TFTP, and HTTP, on the Advance Appliances.
- The appliance terminates TCP connections for incoming DNS requests after handling the initial request through each TCP connection. The exception for this default Grid setting is for an SOA query sent by a client that is accepted in the allow-transfer ACL. In the case of an SOA query, the TCP connection remains open for subsequent DNS requests. This exception also covers the case in which an AXFR query follows the SOA query



through the same TCP connection. For more information about how to override this default Grid setting, see [Enabling Multiple DNS Requests through a Single TCP Session](#).

#### For Software ADP

- When you use IB-FLEX for Software ADP, it supports a standalone or a Grid member with threat protection enabled, but it does not support a Grid Master with threat protection enabled. For more information about the IB-FLEX virtual appliance model, see [About IB-FLEX](#).
- IB-FLEX applies threat protection rules to all traffic on LAN1, LAN2 and HA interfaces, but bypasses the traffic on the MGMT interface.
- The threat protection profiles used for Software ADP members support ADP NAT settings.

On all vNIOS appliances that support vDCA (virtual DNS Cache Acceleration) or vADP (virtual Advanced DNS Protection), you must run vDCA or vADP on a single virtual NUMA node. If the configuration of the virtual NUMA node and physical NUMA node are not the same, it may result in performance degradation.

## Configuring Infoblox Advanced DNS Protection

To enable and configure Infoblox Advanced DNS Protection on supported Infoblox appliances, complete the following:

1. Obtain valid Threat Protection, Threat Protection (Software add-on) and Threat Protection Update licenses from Infoblox and install them on the Infoblox Advance Appliances. For information about license requirements, see [Supported Threat Protection Appliances and Licensing Requirements](#).
2. Optionally, you can set up an HA pair using the same appliance models for both the active and passive nodes. For information about HA pairs, see [About HA Pairs](#). Note that you cannot configure Advance Appliances as an HA Grid Master or HA Grid Master candidate.
3. Enable threat protection service, as described in [Starting and Stopping Threat Protection Service](#). For an HA pair, enable the service on both the active and passive nodes.
4. Configure threat protection rule settings for the Grid, including automatic or manual rule updates, as described in [Configuring Grid Security Properties](#). If your network configuration requires rule updates to go through a proxy server, you can configure the appliance to use a proxy server to send rule updates. For more information, see [Configuring Proxy Servers](#). You can also delegate ruleset updates to specific Grid members and change the default interface to an alternate interface, as described in [Configuring Members and Interfaces for Automatic Updates](#), see [Configuring Proxy Servers](#). Define threat protection profiles for the Grid or specific members, as described in [Configuring Threat Protection Profiles](#).
5. Optionally, you can do the following:
  - Override the default Grid setting that disables multiple DNS requests through one TCP session, as described in [Enabling Multiple DNS Requests through a Single TCP Session](#).
  - Direct incoming DNS queries to be handled by DNS Cache Acceleration before being passed to Threat Protection, as described in [Handling DNS Queries Through DNS Cache Acceleration](#).
  - Modify system rules, as described in [Modifying System and Auto Rules](#).
  - Create custom rules using rule templates, as described in [creating custom rules](#), see [Custom Rules](#).

After you have successfully set up Infoblox Advanced DNS Protection, you can do the following:

- View the current threat protection rules, as described in [Viewing Threat Protection Rules](#).
- Modify system and custom threat protection rules, as described in [Managing Threat Protection Rules](#).
- Manually upload rule updates, as described in [Manually Uploading Rulesets](#).
- Publish uploaded rule updates, as described in [Publishing Rule Updates](#).
- For manual updates, compare differences between two rulesets and merge parameter changes from an old ruleset into a new one, as described in [Comparing and Merging Rulesets](#).
- Monitor threat protection related events and reports, as described in [Monitoring Threat Protection Events](#).
- Set the threat protection service in monitor mode, as described in [Enabling and Disabling Monitoring Mode](#).
- Add threat protection profiles, as described in [Adding Threat Protection Profiles](#).
- Clone threat protection profiles, as described in [Cloning Threat Protection Profiles](#).
- Modify threat protection profiles, as described in [Modifying Threat Protection Profiles](#).
- Merge threat protection profiles, as described in [Merging Threat Protection Profiles](#).
- Inherit Grid rule settings for a threat protection profile, as described in [Inheriting Grid Rule Settings](#).



- Delete a threat protection profile, as described in [Deleting Threat Protection Profiles](#).
- View the current threat protection profiles, as described in [Viewing Threat Protection Profiles](#).

## Administrative Permissions

Superusers can configure all threat protection related tasks. You can assign **Security Permissions** to specific admin groups and roles. You can also add a global permission for managing Grid security properties or add an object permission for managing member security properties.

To perform any action on the rulesets and rules, you must have a Read-write permission for *Grid Security properties* and *Member Security properties* at the Grid level and the member level respectively. For more information about security permissions, see [Administrative Permissions for DNS Threat Protection](#).

## Administrative Permissions for Advanced DNS Protection

Superusers can configure all threat protection related tasks. You can assign **Security Permissions** to specific admin groups and roles. You can also add a global permission for managing Grid security properties or add an object permission for managing member security properties.

To perform any action on the rulesets and rules, you must have a Read-write permission for *Grid Security properties* and *Member Security properties* at the Grid level and the member level respectively. For more information about security permissions, see [Administrative Permissions for DNS Threat Protection](#).

## Supported Threat Protection Appliances and Licensing Requirements

The Infoblox Advanced DNS Protection solution offers the following licenses: **Threat Protection**, **Threat Protection (Software add-on)**, and **Threat Protection Update**. The following are descriptions for each of these licenses:

- **Threat Protection**: Install this license on the physical or hardware-based threat protection appliances. With valid licenses installed, Infoblox Advanced DNS Protection supports both IPv4 and IPv6. You can configure two appliances of the same model to form an HA pair for high availability configuration. For more information about how to configure an HA pair for Infoblox Advanced DNS Protection and its limitations, see [About HA Pairs](#). You can configure hardware-based appliances in either IPv4, IPv6, or dual mode (IPv4 and IPv6) network environment.



### Note

Reverting to an earlier NIOS release that does not support HA configuration for these appliances could cause a service outage.

- **ThreatProtection(Softwareadd-on)**: The Software ADP license is a subscription license. Grid Manager displays a warning message when the license expires. You must renew the license to use the Software ADP service. To renew the license, contact your Infoblox representative or Infoblox Technical Support. Before you obtain your permanent license, you can install a temporary license for **Threat Protection (Software add-on)** using the `set temp_license` CLI command. For more information, refer to the *InfobloxCLIGuide*. Infoblox supports **ThreatProtection(Softwareadd-on)** and **ThreatProtectionUpdate** licenses for elastic scaling. For more information, see [About Elastic Scaling](#).



### Note

You cannot install Multi-Grid Management and Microsoft Management licenses if you install **ThreatProtection** or **ThreatProtection(Softwareadd-on)** licenses on the NIOS appliance.

- **Threat Protection Update**: To receive initial and subsequent threat protection rules and rule updates, you must have the **Threat Protection Update** license installed. You can then configure NIOS to automatically download and publish threat protection rules or you can manually complete the process. For information, see [Manually Uploading Rulesets](#) and [Publishing Rule Updates](#).

With valid licenses installed, the threat protection appliance can be used as a DNS caching or DNS authoritative server. You can join the appliance to the Grid and treat it as a Grid member. Note that if you install a Threat Protection license on a member, you can enable threat protection only on this member. Contact your Infoblox representative to obtain the **Threat Protection**, **Threat Protection Update** and **Threat Protection (Software add-on)** licenses. For information about licenses, see [Managing Licenses](#).

The following table lists all the supported threat protection appliance models and licenses that you can install on them to activate threat protection:

*Threat Protection Appliance Models and Appliances*

Threat Protection Appliance Model	Physical or Virtual	Required License	Comments	
Infoblox-4030-10GE	Physical	Threat Protection and Threat Protection Update	DNS Cache Accelerator Appliance	
PT-1405	Physical		Infoblox Advance Appliances	
PT-2205	Physical			
PT-2205-10GE	Physical			
TE-815	Physical	Threat Protection (Software add-on) and Threat Protection Update	Trinzic appliances	
TE-825	Physical			
TE-1415	Physical			
TE-1425	Physical			
TE-2215	Physical			
TE-2225	Physical			
TE-4015	Physical			
TE-4025	Physical			
TE-V815	Virtual			NIOS virtual appliances
TE-V825	Virtual			
TE-V1415	Virtual			
TE-V1425	Virtual			
TE-V2215	Virtual			
TE-V2220	Virtual			

Threat Protection Appliance Model	Physical or Virtual	Required License	Comments
TE-V2225	Virtual		
TE-V4015	Virtual		
TE-V4025	Virtual		
IB-FLEX	Virtual	Threat Protection (Software add-on) and Threat Protection Update	IB-FLEX

You can install the **Threat Protection (Software add-on)** license on the following hypervisors:

- VMware ESXi with or without SR-IOV enabled
- OpenStack with KVM with or without SR-IOV enabled
- KVM with or without SR-IOV enabled

For more information about these appliances, refer to the respective installation guides under the **Tech Docs** tab on the Infoblox Support site at <https://support.infoblox.com>.

## Starting and Stopping Threat Protection Service

After you install the Threat Protection licenses on the appliance, you can start the threat protection service so you can monitor and mitigate DNS threats on that appliance.

To start or stop threat protection service:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Services** tab, click the Threat Protection service link, and then select the *member* checkbox.
2. From the Toolbar, click **Start** to start the service or **Stop** to stop the service.

Note that when you stop threat protection service, the appliance does not provide visibility or protection against network floods or DNS threats. After you enable threat protection service, you can configure rule settings, add custom rules, and evaluate system rules to ensure that mitigation to DNS threats is handled properly. You can also temporarily disable the threat protection service when necessary. For information about how to configure Grid security settings, see [Configuring Grid Security Properties](#).



### Note

Starting the threat protection service may cause a change in BIND behaviour in handling multiple TCP DNS queries from the same client. To avoid this, you can disable multiple DNS requests by selecting the **Disable multiple DNS requests via single TCP session** checkbox. For more information see, [Enabling Multiple DNS Requests through a Single TCP Session](#).

Note that under normal circumstances when the system is not experiencing any attack, enabling the threat protection service may have a significant performance impact. When the system is under attack, enabling the threat protection service may have a higher performance impact. You might also notice a significant increase in the memory usage due to the threat protection service. You cannot replicate Grid, member, and profile by level threat protection configuration changes during a scheduled full upgrade.

## Enabling and Disabling Monitoring Mode

To ensure that blocking certain DNS queries and packets through threat protection rules does not cause unintended effect on your appliance, you can set the threat protection service in monitor mode. You can also put your appliance in this mode to rule out the possibility that the DNS server is dropping DNS queries. When monitor mode is enabled, the appliance logs DNS packets (instead of dropping them) that would have been blocked by threat protection rules. This

information is recorded in the syslog.

When the Threat Protection service is in monitor mode, the service status changes from **Threat Protection Service is working** to **Threat Protection is working in monitor mode** and the status color changes from green to yellow. The status appears in both the **Data Management** tab -> **Security** tab -> **Members** tab and in the **Grid** tab -> **Grid Manager** tab -> **Services** tab. Note that when one of the members is in monitor mode, the overall status for the Threat Protection service changes from green to yellow. For more information about viewing service status, see [Monitoring Services](#). You can enable or disable the monitor mode for individual Grid members through the CLI command **set smartnic monitor-mode**. You cannot set this configuration at the Grid level. To enable or disable monitor mode for both hardware and Software ADP profiles, you can use the command **set adp monitor-mode on/off**. The **show adp** command displays the status of the monitor mode. Grid Manager displays a warning if the threat protection profile is running in monitor mode. For more information about this command, refer to the *Infoblox CLI Guide*. Note that the **set smartnic monitor-mode** command is recorded in the audit log while the threat protection events are recorded in the syslog. For information about the audit log and syslog, see [Monitoring Tools](#).

## Understanding Threat Protection Rulesets and Rules

To fully implement Infoblox Advanced DNS Protection, ensure that you import the latest threat protection ruleset. To import rulesets, you must have the **Threat Protection Update** license installed on the appliance. For more information, see [Supported Threat Protection Appliances and Licensing Requirements](#). A ruleset comprises all threat protection rules, including system and auto-generated rules, rule templates, custom rules (if any), and parameter definitions and values. For detailed information about threat protection rules, refer to the *Infoblox Threat Protection Rules* available on the Support web site. Infoblox supports a common threat protection ruleset for both hardware and Software ADP members. This ruleset supports all rules and templates. You can also manually upload your rulesets or download rulesets automatically from the IT server.

Infoblox Advanced DNS Protection supports the following threat protection rules:

- Predefined system and auto-generated rules, as described in [System and Auto Rules](#).
- Custom rules, as described in [Custom Rules](#).

Each threat protection rule belongs to a rule category. When you import a ruleset, the appliance publishes the system and auto rules in their respective categories. NIOS automatically manages rule categories and you cannot add, delete, or modify them. It also provides rule templates for creating custom rules. During a ruleset update, some categories and rules may be added or removed. These actions are performed without intervention after the updates are authorized or automatically executed. You cannot add or delete system and auto rules, but you can create custom rules through predefined rule templates and delete them when necessary.

### Note

You can recover only custom rules from the Recycle Bin, if enabled. Rules, rule templates and categories that are removed through ruleset updates are permanently deleted and cannot be restored from the Recycle Bin.

To obtain initial rules and subsequent rule updates, you can configure the appliance to automatically download and publish rulesets or you can manually download them from the Infoblox Support web site and then publish them. For information about how to configure automatic and manual rule settings, see [Configuring Grid Security Properties](#). Note that only the Grid Master receives rules and rule updates. Grid member receives rules and updates through standard Grid replication from the Grid Master. Ruleset data is not replicated to Grid members that do not have the Threat Protection services enabled.

Infoblox recommends that you configure the appliance to automatically receive ruleset updates so your appliance receives the latest rules periodically. If you prefer to manually download and publish rulesets, ensure that you download them frequently to receive the most updated rules. The appliance can store up to nine ruleset versions, and you can select up to five rulesets and switch between these versions for the Grid or members when necessary. For more information about ruleset versions and updates, see [About Ruleset Versions and Updates](#).

## Using the Events Per Second Rule Setting

The **Events per second** setting allows for disabling or throttling of event logs for specific threat protection rules. The default value is one and the maximum value is 700. NIOS displays an error message when you enter a value greater than the maximum value. You can override this event filter at the member level.

Setting the **Events per second** parameter to zero disables logging for that rule. Setting the parameter to any other number enables threat protection logging for that specific rule. For information about how to configure this, see [Configuring Grid Security Properties](#).

Make note of the following guidelines when you enter a value in the **Events per Second per Rule** field:

- The value of this field is applicable only for rules that do not have `event-filter` as part of their format. The following is an example of a rule that has `event-filter` in its format:

```
drop udp any any -> any 53 (msg:"EARLY DROP UDP DNS named author
attempts"; content:"|07|authors|04|bind|00|"; offset:12; sid:110100100;
rev:1;) event_filter gen_id 1, sig_id 110100100, type limit, track by_src,
count 1, seconds 1
```
- For rules that have `event-filter` as part of their format, the event-filter precedes the value in the **Events per Second per Rule** field.
- Therefore, for rules that have `event-filter` as part of their format, to disable event logging for Threat Protection, you must disable it at the rule-level by modifying the syslog file. For rules that do not have `event-filter` as part of their format, to disable event logging for Threat Protection, set 0 in the **Events per Second per Rule** field.

## Custom Rules

Based on your security needs, you can define custom rules using predefined rule templates. Custom rules are typically whitelisting or blacklisting rules. You can create up to 500 custom rules for each rule template offered by Infoblox Advanced DNS Protection . The appliance logs a syslog message if there are more than 500 rules for a specific rule category. You can remove some rules in order to create new ones for that category.



### Note

Before upgrading from an earlier NIOS release to NIOS 6.11.x, ensure that you have less than 500 custom rules for each template. Otherwise, the upgrade may fail.

You can add or delete custom rules for the Grid only. You cannot add or delete them for members, but you can enable, disable, and modify certain rule parameters at the member level.

When you create custom rules, NIOS automatically generates the rule ID from the template used. Note that custom rules do not support IDNs (Internationalized Domain Names). You must first convert IDNs into punycode before entering the data. For information about how to create custom rules, see [Creating Custom Rules](#) below.

When you create custom rules, you are essentially creating whitelisting and blacklisting entries that utilize rate limiting to detect suspicious UDP and TCP traffic. Infoblox Advanced DNS Protection supports a series of rule templates for defining new custom rules. For information about rule templates, see the following [Custom Rule Templates](#) section. Whitelisting rules define a list of allowed resources before they are blocked by the configured rate limit settings. They provide for only a selected set of entities to access the protected environment. Examples include company offices and their associated internal network services, which presumably use access control systems to enforce them. In effect, addresses or networks that do not match the whitelisting entries are automatically blocked.

Blacklisting rules define a list of disallowed resources through FQDN lookups as well as rate limiting. Blacklists typically allow a far broader base of access to many more entities, and cite a list of specific entities or people that do not have access. Otherwise, any devices or users theoretically have access to the protected environment.

For whitelist entries, the matching values are mandatory, in which the IP address or network of the rule is expressly

permitted access. Blacklist entries are forbidden, in which the IP address or network of the rule is expressly denied access. In essence, blacklisting is convenience while whitelisting is more secure.

 **Note**

You can create DNS-specific blacklists under the **Data Management** tab → **DNS** tab. However, you cannot use this blacklist feature as part of DNS threat protection.

## Custom Rule Templates

Infoblox Advanced DNS Protection supports a few custom rule templates from which you create new custom rules. Note that when you use a specific rule template to create custom rules, the new rules reside in their respective rule categories. For information about creating custom rules, see [Creating Custom Rules](#) below.

For each rule you create, you can define the **Events per second** value to determine the number of events per second that will be logged for the rule. You can also define specific rule parameters for custom rules, as follows:

When you create custom rules that involve FQDN lookups, the appliance automatically verifies the FQDN syntax and format that you enter in the **Value** field. It properly translates escaped sequences and special characters that are used to represent specific characters in the FQDN. For example, \032 is interpreted as a space (hex 20), and \" is interpreted as the double quote (hex 22). The appliance sends an error message when it detects invalid characters in the FQDN.

 **Note**

Custom rules do not support IDNs (Internationalized Domain Names). To use IDNs for custom rules, you must first convert the IDNs into punycode. You can use the **IDN Converter** from the **Toolbar** for the conversion.

- **BLACKLIST FQDN lookup TCP:** Use this rule template to create custom rules for blacklisting DNS queries by FQDN lookups on TCP. In the Rule Parameters table, complete the following:
  - **Blacklisted FQDN:** Enter the FQDN that you want the appliance to block over TCP traffic. NIOS supports an exact match or subdomain matches for the FQDN specified in the rule. For example, if "test.com" is specified as a custom rule, NIOS blocks "test.com" or "abc.test.com" but "abctest.com" will not be blocked.
- **BLACKLIST FQDN lookup UDP:** Use this rule template to create custom rules for blacklisting DNS queries by FQDN lookups on UDP. In the Rule Parameters table, complete the following:
  - **Blacklisted FQDN:** Enter the FQDN that you want the appliance to block over UDP traffic. NIOS supports an exact match or subdomain matches for the FQDN specified in the rule. For example, if "test.com" is specified as a custom rule, NIOS blocks "test.com" or "abc.test.com" but "abctest.com" will not be blocked.
- **BLACKLIST IP TCP Drop prior to rate limiting:** Use this rule template to create rules for blocking IPv4 or IPv6 addresses on TCP before the appliance drops the packets based on rate limiting rules you have defined using the **BLACKLIST IP TCP Drop prior to rate limiting** template. In the Rule Parameters table, complete the following:
  - **Blacklisted IP address/network:** Enter the IPv4 or IPv6 address from which packets sent are dropped before any relevant rate limiting rules take effect. Note that all TCP traffic from the specified IPv4 and IPv6 addresses and networks will be blocked. Enter network addresses in address/CIDR format.
- **BLACKLIST IP UDP Drop prior to rate limiting:** Use this rule template to create rules for blocking IPv4 or IPv6 addresses on UDP before the appliance drops the packets based on rate limiting rules you have defined using the **BLACKLIST IP UDP Drop prior to rate limiting** template. In the Rule Parameters table, complete the following:
  - **Blacklisted IP address/network:** Enter the IPv4 or IPv6 address from which packets sent are dropped before any relevant rate limiting rules take effect. Note that all UDP traffic from the specified IPv4 and IPv6 addresses and networks will be blocked. Enter network addresses in address/CIDR format.
- **RATELIMITED FQDN lookup TCP:** Use this template to create custom rules that contains rate limiting restrictions for blocking DNS queries by FQDN lookups on TCP traffic. In the Rule Parameters table, complete the following:
  - **Packets per second:** Enter the number of packets per second to define the rate limit for this rule. You define this value to control the rate of TCP traffic that consists of DNS lookups for the FQDN defined in this rule. The default is 5.
  - **Drop interval:** Enter the number of seconds for which the appliance drops packets.
  - **Blacklist rate limited FQDN:** Enter the FQDN that is affected by the rate limit value configured for this rule. The appliance drops the packets sent by this FQDN when the TCP traffic of DNS lookups for this FQDN exceeds the configured rate limit value.



- **RATELIMITED FQDN lookup UDP:** Use this rule template to create custom rules that contains rate limiting restrictions for blocking DNS queries by FQDN lookups on UDP traffic. In the Rule Parameters table, complete the following:
  - **Packets per second:** Enter the number of packets per second to define the rate limit for this rule. You define this value to control the rate of UDP traffic that consists of DNS lookups for the FQDN defined in this rule. The default is 5.
  - **Drop interval:** Enter the number of seconds for which the appliance drops packets.
  - **Blacklist rate limited FQDN:** Enter the FQDN that is affected by the rate limit value configured for this rule. The appliance drops the packets sent by this FQDN when the UDP traffic of DNS lookups for this FQDN exceeds the configured rate limit value.

 **Note**

Make sure that you enter a valid FQDN. Example: [test.com](#), [foo.com](#), etc. The appliance does not display an error message when you enter an invalid FQDN. However, the Threat Protection dashboard displays a warning message for the invalid FQDNs. For information about threat protection status for Grid, see [Status Dashboard](#).

- **RATELIMITED IP TCP:** Use this rule template to create custom rules that contains rate limiting restrictions for blacklisting IP addresses on TCP. If there are certain IP addresses that you want to block before its traffic reaches the rate limit restrictions, you can create a rule using the **RATELIMITED IP TCP** template. In the Rule Parameters table, complete the following:
  - **Packets per second:** Enter the number of packets per second to define the rate limit for this rule. You define this value to control the rate of TCP traffic that consists of DNS lookups for the IP address or network defined in this rule. The default is 5.
  - **Drop interval:** Enter the time interval in seconds the appliance drops IP packets sent by the rate limited IP address or network defined for this rule. The default is 30 seconds.
  - **Rate limited IP address/network:** Enter the IP address or network that is affected by the rate limit value configured for this rule. The appliance drops the packets sent by this IP address based on the drop interval when the TCP traffic of DNS lookups for this IP address exceeds the configured rate limit value.
- **RATELIMITED IP UDP:** Use this rule template to create custom rules that contains rate limiting restrictions for blacklisting IP addresses on UDP. If there are certain IP addresses that you want to block before its traffic reaches the rate limit restrictions, you can create a rule using the **RATELIMITED IP UDP** template. In the Rule Parameters table, complete the following:
  - **Packets per second:** Enter the number of packets per second to define the rate limit for this rule. You define this value to control the rate of UDP traffic that consists of DNS lookups for the IP address or network defined in this rule. The default is 5.
  - **Drop interval:** Enter the time interval in seconds the appliance drops IP packets sent by the rate limited IP address or network defined for this rule. The default is 30 seconds.
  - **Rate limited IP address/network:** Enter the IP address or network that is affected by the rate limit value configured for this rule. The appliance drops the packets sent by this IP address based on the drop interval when the TCP traffic of DNS lookups for this IP address exceeds the configured rate limit value.
- **WHITELIST IP TCP Pass prior to rate limiting:** Use this rule template to create custom rules for allowing certain IP addresses on TCP before the appliance drops the packets based on rate limiting rules you have defined using the **RATELIMITED IP TCP** template. In the Rule Parameters table, complete the following:
  - **Whitelisted IP address/network:** Enter the IPv4 or IPv6 address from which packets sent are allowed before any relevant rate limiting rules take effect.
- **WHITELIST IP UDP Pass prior to rate limiting:** Use this rule template to create custom rules for allowing certain IP addresses on UDP before the appliance drops the packets based on rate limiting rules you have defined using the **RATELIMITED IP UDP** template. In the Rule Parameters table, complete the following:
  - **Whitelisted IP address/network:** Enter the IPv4 or IPv6 address from which packets sent are allowed before any relevant rate limiting rules take effect.
- **WHITELIST TCP Domain:** Use this rule template to create custom rules to allow DNS queries by FQDN lookups on TCP. In the Rule Parameters table, complete the following:
  - **Whitelist FQDN:** Enter the FQDN that you want the appliance to allow over TCP traffic. NIOS supports an exact match or subdomain matches for the FQDN specified in the rule. For example, if "test.com" is specified as a custom rule, NIOS blocks "test.com" or "abc.test.com" but "abctest.com" will not be blocked.

- **WHITELIST UDP Domain:** Use this rule template to create custom rules to allow DNS queries by FQDN lookups on UDP. In the Rule Parameters table, complete the following:
  - **Whitelist FQDN:** Enter the FQDN that you want the appliance to allow over UDP traffic. NIOS supports an exact match or subdomain matches for the FQDN specified in the rule. For example, if "test.com" is specified as a custom rule, NIOS blocks "test.com" or "abc.test.com" but "abctest.com" will not be blocked.
- **BLACKLIST TCP FQDN lookup for DNS Message Type:** Use this rule template to create custom rules for blacklisting FQDN lookups on TCP for the specified DNS message type. In the Rule Parameters table, complete the following:
  - **DNS Record Type:** Select the DNS record type from the drop-down list or enter a valid ENUM for the DNS record. You can enter a value between 1 and 65534. The following DNS resource records are not supported by this rule template: MD (3), MF (4), MB (7), MG (8), MR (9), WKS (11), HINFO (13), MINFO (14), IXFR (251), and AXFR (252) record.
  - **Blacklisted FQDN substring:** Enter the FQDN from which the packets received are blocked over TCP for the specified DNS message type.
- **BLACKLIST UDP FQDN lookup for DNS Message Type:** Use this rule template to create custom rules for blacklisting FQDN lookups on UDP for the specified DNS message type. In the Rule Parameters table, complete the following:
  - **DNS Record Type:** Select the DNS record type from the drop-down list or enter a valid ENUM for the DNS record. You can enter a value between 1 and 65534. The following DNS resource records are not supported by this rule template: MD (3), MF (4), MB (7), MG (8), MR (9), WKS (11), HINFO (13), MINFO (14), IXFR (251), and AXFR (252) record.
  - **Blacklisted FQDN substring:** Enter the FQDN from which the packets received are blocked over UDP for the specified DNS message type.
- **Pass TCP DNS Message Types:** Use this rule template to create custom rules to allow TCP DNS packets that contain the specified DNS record type. In the Rule Parameters table, complete the following:
  - **DNS Record Type:** Select the DNS record type from the drop-down list or enter a valid ENUM for the DNS record. You can enter a value between 1 and 65534. The following DNS resource records are not supported by this rule template: MD (3), MF (4), MB (7), MG (8), MR (9), WKS (11), HINFO (13), MINFO (14), IXFR (251), and AXFR (252) record.
- **Pass UDP DNS Message Types:** Use this rule template to create custom rules to allow UDP DNS packets that contain the specified DNS record type. In the Rule Parameters table, complete the following:
  - **DNS Record Type:** Select the DNS record type from the drop-down list or enter a valid ENUM for the DNS record. You can enter a value between 1 and 65534. The following DNS resource records are not supported by this rule template: MD (3), MF (4), MB (7), MG (8), MR (9), WKS (11), HINFO (13), MINFO (14), IXFR (251), and AXFR (252) record.
- **RATE LIMITED TCP DNS Message Type:** Use this rule template to create custom rules that contain rate limiting restrictions for blacklisting TCP DNS packets that contain the specified DNS record type. In the Rule Parameters table, complete the following:
  - **Packets per second:** Enter the number of packets per second to define the rate limit for this rule. You define this value to control the rate of TCP traffic that consists of DNS packets with the DNS record type defined in this rule. The default is 5.
  - **DNS Record Type:** Select the DNS record type from the drop-down list or enter a valid ENUM for the DNS record. You can enter a value between 1 and 65534. The following DNS resource records are not supported by this rule template: MD (3), MF (4), MB (7), MG (8), MR (9), WKS (11), HINFO (13), MINFO (14), IXFR (251), and AXFR (252) record.
  - **Drop interval:** Enter the number of seconds for which the appliance drops packets.
- **RATE LIMITED UDP DNS Message Type:** Use this rule template to create custom rules that contain rate limiting restrictions for blacklisting UDP DNS packets that contain the specified DNS record type. In the Rule Parameters table, complete the following:
  - **Packets per second:** Enter the number of packets per second to define the rate limit for this rule. You define this value to control the rate of UDP traffic that consists of DNS packets with the DNS record type defined in this rule. The default is 5.
  - **DNS Record Type:** Select the DNS record type from the drop-down list or enter a valid ENUM for the DNS record. You can enter a value between 1 and 65534. The following DNS resource records are not supported by this rule template: MD (3), MF (4), MB (7), MG (8), MR (9), WKS (11), HINFO (13), MINFO (14), IXFR (251), and AXFR (252) record.
  - **Drop interval:** Enter the number of seconds for which the appliance drops packets.



## Creating Custom Rules

Infoblox Advanced DNS Protection provides a few rule templates from which you can create custom rules. For information about the list of rule templates that you can use, see [Custom Rule Templates](#).

To create a custom rule:

1. From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab -> *Ruleset* link, and then click **Add Custom Rule** from the Toolbar.
2. In the *Add Custom Rule* editor, complete the following:
  - **Template:** From the drop-down list, select the blacklisting or whitelisting rule template from which you want to create the new rule. For more information about the rule templates, see the previous section, Custom Rule Templates.
  - **Description:** Displays the description of the rule that you are about to create. You cannot modify this.
  - **Comment:** Enter comments to describe the new rule.
  - **Disable:** Select this if you want to keep the new rule disabled for later use.
3. Click **Next** and complete the following to configure rule parameters:
  - **Description:** Displays the description of the rule that you are about to create. You cannot modify this.
  - **Action:** Displays the operation the appliance performs when an event related to this rule occurs. Some rules are restricted to specific actions. For example, the action for all blacklisting rules is set to **Drop**, where the appliance drops IP packets when such an event occurs. The action for all whitelisting rules is set to **Pass**, where the appliance passes IP packets when such an event occurs.
  - **Log Severity:** Select **Critical**, **Major**, **Warning** or **Informational**. The log severity you select here determine the severity of the message triggered by a match against the rule.

In the Rule Parameters section, do the following:

- Click the **Value** field for **Events per second** to enter the maximum number of events per second that will be logged for this rule. Use a nonzero value if you want matches against the current rule to log events. Setting a value to 0 (zero) disables the appliance from logging events associated with this rule.
  - Depending on the template you have selected, click the **Value** field and enter the appropriate parameters to configure the rule. For descriptions about the parameters for each rule template and System and Auto Rule Categories, see [System and Auto Rules](#). Note that when you create custom rules that involve FQDN lookups, the appliance automatically verifies the FQDN syntax and format that you enter in the **Value** field. It properly translates escaped sequences and special characters that are used to represent specific characters in the FQDN. For example, \32 is interpreted as a space (hex 20), and \" is interpreted as the double quote (hex 22). The appliance sends an error message when it detects invalid characters in the FQDN.
4. Click **Save & Close**.

The new rule, with an automatically assigned rule ID, is created and placed in its corresponding rule category.

## System and Auto Rules

System rules are predefined threat protection rules that are built into the Advance Appliances. New system rules are added through rule updates. You can enable an entire category of system rules as well as individual rules. Note that you cannot add or delete system rules, though you can change some parameters. For most system rules, you can modify the **Action** and **Log Severity**. For more information, see [Modifying System and Auto Rules](#).

Auto-generated rules are firewall rules that are automatically defined by NIOS for blocking traffic for disabled services and ports. They do not support functionality such as rate limiting. These rules can be grouped into different rule categories and are enabled or disabled by default. You cannot enable or disable auto rules in this release of Infoblox Advanced DNS Protection, though you can set log severity and control logging for these rules. Note that auto rules are automatically enabled or disabled and reconfigured based on the current running services and configuration on the appliance.

## System and Auto Rule Categories

The appliance supports the following system and auto rule categories. For detailed descriptions about each system and auto rule, refer to the Threat Protection Rules document available on the Infoblox Support site.

- **BGP:** Contains auto rules that mitigate attacks that target BGP (Border Gateway Protocol) routing parameters, such as invalid attribute lengths or invalid message types.
- **DNS Amplification and Reflection:** Contains system and auto rules that can be used to mitigate the commonly used methods of DDoS attacks. For information about DNS amplification and reflection, see [DNS Reflection and Amplification Attacks](#).
- **DNS Cache Poisoning:** Contains rate limiting rules that assign bandwidth restrictions rules used to mitigate DNS cache poisoning (on UDP and TCP) that is performed by sending a large volume of fake replies to a recursive server, which can result in hundreds or thousands of redirects. For more information about DNS cache poisoning, see [DNS Cache Poisoning](#).
- **DNS Malware:** Contains rules that protect against DNS malware that posts serious threats to the DNS infrastructure. For information about DNS malware threats, see [DNS Malware](#).
- **DNS DDoS:** Contains system rules that are used to mitigate DNS DDoS attacks on your Advance Appliance. These rules rate limits clients that trigger following DNS responses: NXDOMAIN, NXRRSET, and SERVFAIL.
- **DNS Message Type:** Contains DNS system rules that can be used to filter requests that query specific DNS flags in the DNS message header.
- **DNS Protocol Anomalies:** Contains auto rules that address general DNS protocol attacks such as invalid DNS queries.
- **DNS Tunneling:** Contains auto rules that mitigate against DNS tunneling attacks. For more information, see [Inside-Out Attacks](#).
- **Default Drop:** Contains system rules that automatically drop IP packets when unusual UDP, TCP, and ICMP traffic is detected.
- **General DDoS:** Contains auto rules that address general DDoS (Distributed Denial of Service) attacks such as loopback address spoofing, and UDP or TCP packets that contain the same source and destination addresses.
- **HA Support:** Contains auto rules that are used to pass packets that go through the Virtual Router Redundancy Protocol (VRRP) and Internet Group Management Protocol (IGMP) for HA (High Availability) support.
- **ICMP:** Contains auto rules that mitigate ICMP and ICMPv6 ping attacks. ICMP ping size (for IPv4 and IPv6) for these rules is limited to 792 bytes. For information about ICMP, see [Internet Control Message Protocol \(ICMP\) Flood](#).
- **NTP:** Contains auto rules that mitigate attacks that target the NTP (Network Time Protocol). These rules include support for NTP requests and responses, NTP IPv4 and IPv6 ACLs (Access Control Lists), and private mode 7 packets.
- **OSPF:** Contains auto rules that mitigate attacks that target OSPF (Open Shortest Path First) routing parameters, such as invalid attribute lengths or invalid message types.
- **Potential DDoS Related Domains:** Contains system rules the appliance uses to blacklist domains that may have been the targets or subjects in NXDOMAIN or DDoS attacks. These rules block all FQDN lookups on UDP for domains that have been observed to be used as targets in DDoS attacks, and they are enabled by default.
- **Reconnaissance:** Contains auto rules that mitigate network reconnaissance attacks, in which unauthorized remote attackers attempt to access networks by exploiting network standards and communications.
- **TCP/UDP Floods:** Contains DNS system rules that are used to mitigate DNS TCP and UDP floods. For information about TCP/UDP floods, see [UDP DNS Flood](#).

## About Ruleset Versions and Updates

Infoblox periodically releases updated rulesets. After an automatic update, the new ruleset is automatically applied to the Grid members that are using the Grid wide ruleset. After a manual update, you can manually apply the new ruleset to the Grid or individual Grid members. Before you manually publish a ruleset, you can view differences between the current ruleset and the newly downloaded one. You can also modify some changed parameters and then merge the changes from the old version to the new one. For more information, see [Comparing and Merging Rulesets](#). For information on applying rulesets based on profiles, see [Adding Threat Protection Profiles](#).

You can view the current ruleset version (displayed in the **Version** column) in one of the following tabs:

- **Grid:** From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab.
- **Member:** From the **Data Management** tab, select the **Security** tab -> **Members** tab.
- **Profile:** From the **Data Management** tab, select the **Security** tab -> **Profiles** tab.

In a Grid, you can run different versions of rulesets on different Grid members. For example, Grid member 1 can use revision 1 and Grid member 2 can use revision 2 of the ruleset. You can also switch back to a previous ruleset version when necessary, but you cannot change the version number for individual rules.

The appliance retains up to nine (9) rulesets at any given time: five (5) old rulesets, one (1) newly downloaded ruleset

and three (3) "Do Not Delete" rulesets. You can configure rulesets as "Do Not Delete" at the Grid level. The appliance retains these rulesets and they cannot be deleted during an automatic or manual ruleset update. To allow a ruleset to be deleted during an update, you must first disable the "Do Not Delete" flag for these rulesets. Note that you cannot delete a ruleset that is used by the Grid or any members. For information about how to configure the "Do Not Delete" flag, see [Modifying Rulesets](#).

For more information about how to add, modify, and delete a ruleset, see [Managing Threat Protection Rulesets](#).

## Ruleset Update Behavior

Consider the following behavior during a ruleset update:

- If you have configured the rule update policy as **Automatic**, the following occurs:
  - For each rule that exists in the current ruleset and is used by a Grid member, the appliance automatically copies all customized parameter values from the current ruleset to the corresponding rules in the new ruleset. For rule templates that exist in the current ruleset and are used by any member in the Grid, the appliance automatically copies the rule instances from the current version to the new ruleset.
  - The appliance automatically compares and integrates all rule changes into the new ruleset. For manual ruleset update, you can view the differences between the current ruleset and the newly downloaded ruleset, select specific rules and make modifications to customized parameters, and then merge the changes into the new ruleset before applying it to the Grid and members. For more information about how to view the rule differences and merge changes, see [Comparing and Merging Rulesets](#).
  - The new ruleset is applied to the Grid and only members that use the same ruleset version as the Grid through inheritance from the Grid security properties. For members that use the same ruleset version as the Grid but has parameter overrides for certain rules, the overridden values will be copied to the new ruleset.

For information about how to configure the rule update policy, see [Configuring Grid Security Properties](#).

- When there are a total of nine (9) rulesets stored in the database, the ruleset that is not used by the appliance and is not marked as "Do Not Delete" will be replaced by the newly downloaded ruleset.
- If there are more than one ruleset that can be replaced, the appliance selects the oldest version based on the version number.
- If a specific system or auto rule from the current ruleset version does not exist in the new ruleset, it will not be migrated to the new ruleset.
- If a specific template from the current ruleset version does not exist in the new ruleset, all of its custom rules will not be migrated to the new ruleset.

## Packet Flow for Threat Protection Rules

Threat protection rules are designed to work together to provide maximum protection for your environment. This section describes how these rules are being applied and how you can tune some of them to suit your system setup and network environment.

Threat protection rules are grouped by rule categories, and most of them have one or more associated rule parameters. All threat protection rules contain rule parameters that you may or may not be able to configure. Rule parameters are predefined with default values that generally suit most network environments. However, there are times when you have special setups or configurations in your environment that require special attention. In these cases, you may need to change some of the rule parameters to obtain optimal protection without sacrificing system performance. For detailed information about all threat protection rules and how to tune them when necessary, refer to the *Infoblox Threat Protection Rules* available on the Support web site.

Depending on the rules, you may or may not be able to override default values for the following rule parameters (when applicable):

- **Packets per second:** This parameter defines the rate limit or the number of packets per second that the appliance processes before it performs a triggered action, such as sending warnings or blocking traffic.
- **Drop interval:** This is the time period (in seconds) for which the appliance blocks traffic from the client or traffic that matches a certain pattern beyond the rate limit. Based on how you want to handle the traffic that exceeds the rate limit, you can configure this interval to work with the **Rate Algorithm** parameter.
- **Rate algorithm:** This parameter defines how the appliance handles incoming traffic when the traffic exceeds the rate limit (defined in **Packets per second**). You can set this to "blocking" or "rate limiting." The default is "rate limiting." When you set this to "blocking," the appliance allows client traffic to go through until it hits the rate limit.

It then blocks all traffic for the duration of the drop interval. If client traffic continuously exceeds the rate limit, the appliance continues to block all traffic for subsequent drop intervals without letting through any traffic, which could result in an indefinite traffic blockage. When you set this to "rate limiting," the appliance allows client traffic to go through until traffic hits the rate limit. It then blocks all traffic for the rest of the drop interval. The appliance re-evaluates client traffic at the beginning of each drop interval and repeats the same behavior for subsequent intervals.

To avoid resource exhaustion and limit frauds, you can limit the query rate for each source IP, and then set **Dropinterval** to one second and **Ratealgorithm** to "rate limiting," which results in a rate-limiting behavior that allows some traffic to go through before the rest of the traffic is blocked. In this case, the appliance re-evaluates the client behavior every second. If the client traffic exceeds the rate limit, the appliance processes only queries up to the rate limit and drops all excessive queries for the remainder of the second.

For more information about how to configure **Ratealgorithm**, **Packetspersecond** and **Dropinterval**, see the following section, Configuration Examples.



#### Note

Starting with NIOS 6.12.4, the default for **Ratealgorithm** has been changed from "blocking" to "rate limiting."

- **Events per second**: The number of events logged per second for the rule. Setting a value to 0 (zero) disables the appliance from logging events for the rule. Most rules have this parameter, and the default value is 1.
- **Packet size**: DNS packet size. If the DNS packet size exceeds a certain value, the corresponding rule will be triggered.

### Configuration Examples

Depending on how you want the appliance to handle incoming traffic, you can configure applicable parameters so they work hand-in-hand to deliver desired results. Following are some examples that demonstrate how you can use the **Rate algorithm**, **Packets per second** and **Drop interval** parameters.

#### Example 1

If source IP 100.10.10.1 sends queries at a rate of 100 packets per second, and you have the following configuration for a threat protection rule:

**Packets per second** = 40

**Drop interval** = 3

**Rate algorithm** = blocking

The appliance handles incoming traffic in the following manner:

1st second: 40 packets are allowed; all other packets are blocked

2nd second: All traffic from 100.10.10.1 is blocked

3rd second: All traffic from 100.10.10.1 is blocked

4th second: All traffic from 100.10.10.1 is blocked

5th second: All traffic from 100.10.10.1 is blocked

6th second: All traffic from 100.10.10.1 is blocked

In this example, the appliance evaluates historic data of the client behavior. If the client traffic exceeds the limit, then the appliance continuously drops traffic, which may result in an indefinite traffic blockage for a client that continuously violates the rate limit.

#### Example 2

Source IP 100.10.10.1 sends queries at a rate of 100 packets per second for a duration of two seconds. It then sends 30 packets per second for three seconds and 50 packets afterwards; and you have the following configuration for a threat protection rule:

**Packets per second** = 40

**Drop interval** = 1

**Rate algorithm** = rate limiting

The appliance handles incoming traffic in the following manner:

1st second: 40 packets are allowed; all other packets are blocked for the remainder of the second

2nd second: 40 packets are allowed; all other packets are blocked for the remainder of the second

3rd second: All traffic from 100.10.10.1 is allowed

4th second: All traffic from 100.10.10.1 is allowed

5th second: All traffic from 100.10.10.1 is allowed

6th second: 40 packets are allowed; all other packets are blocked for the remainder of the second

In this case, the appliance re-evaluates the client behavior every second. If client traffic exceeds the rate limit, the appliance processes queries up to the rate limit and drops all excessive queries for the remainder of the second.

## Configuring Grid Security Properties

After you have installed valid threat protection licenses, you can configure rule update settings for the Grid. The Grid settings apply to all members in the Grid. You can select to use an existing threat protection ruleset or use a threat protection profile. A threat protection profile defines specific security settings and ruleset that you want to apply to a specific member or a group of members. For more information about rulesets, see [Understanding Threat Protection Rulesets and Rules](#). For information about threat protection profiles, see [Configuring Threat Protection Profiles](#). You can override only the global **Event per second** filter (in the **Basic** tab) and the **Disable multiple DNS requests via single TCP session** option (in the **Advanced** tab) in the *Member Security Properties* editor by selecting a member and clicking Edit.

To configure rule settings for the Grid or an individual member:

1. **Grid:** From the **Data Management** tab, select the **Security** tab, and then click **Grid Security Properties** from the Toolbar.
2. In the *Grid Security Properties* editor, select the **Threat Protection** tab -> **Basic** tab, and complete the following:
  - **Active Ruleset Version:** To activate and apply a specific ruleset to the Grid, select the ruleset from the drop-down list.
  - **Active Ruleset Comment:** Displays information about the selected ruleset from the **Comment** field of the ruleset.  
In the *Threat Protection Ruleset Updates* section, define the rule update policy. The appliance automatically performs rule updates by default. You can choose to manually publish rule updates. For information about how to manually update rules, see [Manually Uploading Rulesets](#) and [Publishing Rule Updates](#).
  - **Latest Available Ruleset:** Displays the version string of the last published ruleset. This field changes each time when the ruleset is updated.
  - **Last Checked For Updates:** Displays the timestamp and time zone when you manually upload a ruleset file or automatically download the latest rule update file from the Infoblox rule update server. This field changes the timestamp only when there is a change in the ruleset. The appliance does not update this field if there is no change in the ruleset for the manual upload or automatic download.
  - **Rule Update Policy:** Select the rule update policy from the drop-down list to determine whether updates are being applied automatically or manually. When you select **Automatic**, the appliance automatically switches to the newly downloaded ruleset and publishes the changes when a ruleset update is applied. This schedule is recurrent and occurs every 24 hours. Note that the download happens either 15 minutes before or after the time specified. Select **Manual** to manually download updated rulesets and publish them. Note that you must have a valid Threat Protection Update license installed in order to perform ruleset updates. For information about how to perform a manual update, see [Manually Uploading Rulesets](#) and [Publishing Rule Updates](#).

- **Enable Automatic Ruleset Downloads:** Select this to enable automatic ruleset downloads. Note that starting with NIOS 8.0.0, ruleset downloads might take longer than previous releases; but there is no functional impact during the downloads.



#### Note

When you select this, ensure that you configure and enable a valid DNS resolver for the Grid in the *Grid Properties* editor so the appliance can successfully access the updated ruleset file.

If your network environment does not allow direct HTTP or HTTPS communication with the Internet through a firewall from a secure location in which the Grid Master or the standalone appliance resides, you can configure the Advance Appliance to use a proxy server so you can receive automatic threat protection updates through this connection. Configured proxy settings are for the entire Grid. You cannot configure proxy settings for individual members. For information about how to configure proxy servers, see [Configuring Proxy Servers](#).

- **Test Connection:** Click this to test the connectivity between the Advance Appliance and the server from which you receive the rule update files. Grid Manager displays a message indicating whether the connection is successful.
  - **Download Rules Now:** Click this to immediately download the latest rule update file from the Infoblox rule update server, provided that the connection between the appliance and the server is successful.
- In the Schedule section, define the schedule for automatic ruleset downloads. The following options are enabled only when you have selected **Enable Automatic Ruleset Downloads**:
- **Default:** Select this to set the default schedule settings for automatic ruleset downloads. When you select the default schedule, download starts anywhere between 12 AM and 6 AM local time.
  - **Custom:** Select this to schedule downloads at a later date and time. Click the Calendar icon to select the date and time.



#### Note

When you schedule automatic ruleset downloads, the downloads are performed within 15 minutes before or after the scheduled time. If you have multiple Grid members configured for downloads, the same offset time applies to all members when the first member is unreachable. Downloads to the next reachable member do not happen right after a download fails on the unreachable member. The offset time is put in place to prevent all members from performing downloads at the same time.

In the Threat Protection Logging section, define the events per second per rule value to allow the appliance to log events in the syslog:

- **Events per Second per Rule:** Specify the number of events logged per second per rule. The default value is one and the maximum value is 700. Setting the value to 0 (zero) disables NIOS from logging events for the rules. NIOS displays an error message when you enter a value greater than the maximum value. You can override this event filter at the member level. For guidelines about using this setting, see [Using the Events Per Second Rule Setting](#).

3. To save the configuration, click **Save & Close**. To publish changes, click **Publish** if it appears at the top of the screen. Note that NIOS does not require restarting of the threat protection service after rule updates.



#### Note

When you enable the **Enable Automatic Ruleset Downloads** option and set the **Rule Update Policy** to Manual, and download a new ruleset manually, you need to click **Save & Close** to retain the existing customized rules. If you click **Cancel** or close the wizard, the existing customized ruleset will be replaced by the downloaded ruleset.



## Enabling Multiple DNS Requests through a Single TCP Session

The advanced appliance inspects only one DNS request sent over a single TCP connection. To avoid accepting possible malicious data following a valid DNS request, the appliance terminates the TCP connection after handling the initial DNS request over TCP. You can modify this default Grid setting at the Grid or member level.

To modify this setting, do the following:

1. **Grid:** From the **Data Management** tab, select the **Security** tab, and then click **Grid Security Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **Security** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.  
**Profiles:** From the **Data Management** tab, select the **Security** tab -> **Profiles** tab -> *profile* checkbox, and then click the Edit icon.
2. In the *Grid Security Properties* or *Member Security Properties* editor, select the **Threat Protection** tab -> **Advanced** tab, and complete the following:
  - **Disable multiple DNS requests via single TCP session:** This is selected by default to avoid accepting possible malicious data following a valid DNS request. When this is selected, the appliance handles the initial DNS request through TCP and then terminates the TCP session to block subsequent DNS traffic, except for an SOA query sent by a client that is accepted in the allow-transfer ACL. This exception covers the case in which an AXFR query follows the SOA query through the same TCP connection. This field is read-only when you use a threat protection profile instead of a ruleset. For more information, see [Configuring Grid Security Properties](#).
3. Save the configuration.



### Note

The **Disable multiple DNS requests via single TCP session** checkbox is enabled by default when Advanced DNS Protection is enabled.

## Handling DNS Queries Through DNS Cache Acceleration

By default, Threat Protection rules are applied on all DNS queries and packets before they are passed on to DNS Cache Acceleration. However, you can configure such that DNS queries and packets are first passed on to DNS Cache Acceleration. If the query is valid and the answer is in the cache, the query is answered by DNS Cache Acceleration. Only if the answer is not in the cache, Threat Protection rules are applied to the query.

Passing DNS queries and packets to DNS Cache Acceleration first better the performance for recursive DNS queries.

To handle DNS Queries Through DNS Cache Acceleration, do the following:

1. **Grid:** From the **Data Management** tab, select the **Security** tab, and then click **Grid Security Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **Security** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.
2. In the *Grid Security Properties* or *Member Security Properties* editor, select the **Threat Protection** tab -> **Advanced** tab, and complete the following:  
**Enable DNS responses from acceleration cache before applying Threat Protection rules:** Select this checkbox for incoming DNS queries to be first handled by DNS Cache Acceleration. This checkbox is disabled by default.
3. Save the configuration.

## Guidelines for Enabling DNS Responses from DNS Cache Acceleration

The following guidelines apply when you select the **Enable DNS responses from acceleration cache before applying Threat Protection rules** checkbox:

- This checkbox is available only on IB-FLEX platforms which are capable of both Advanced DNS Protection and DNS Cache Acceleration.
- You do not need to restart the DNS service or publish Threat Protection rules for the setting to take effect.

## Limitation of Enabling DNS Responses from DNS Cache Acceleration

If you select the **Enable DNS responses from acceleration cache before applying Threat Protection rules** checkbox, rate-limiting or response based rules will not be triggered if the FQDNs queries are being cached.

## Configuring NAT Mapping Properties

If you are operating in an environment where you need to aggregate multiple users (such as mobile users) behind one IP address, you might want to consider using NAT (Network Address Translation) Mapping to avoid the potential of service denial to a large group of users if threat protection rules are being applied to the same IP address.

NAT detection is essential to enforce the correct threat protection rule on source IP addresses. The Infoblox NAT Mapping feature allows you to designate individual IP addresses, networks, or ranges along with the source port ranges to denote NAT'ed clients. Any UDP or TCP traffic that originates from an IP address within the NAT IP range and uses a source port within the NAT port mapping is considered as a NAT'ed client. It is important to align the NIOS NAT mapping configuration with the NAT gateway configuration to avoid mis-classification of NAT IPs. Traffic that originates from a given source port block that falls within the configured NAT IP and port range is considered as the same NAT'ed client and traffic that originates from a different port block that falls within the same NAT IP and port range is considered to be a different NAT'ed client. If the traffic originates from the same source IP address, but from a different port block that falls outside the configured port range, then they are considered as a non-NAT'ed client.

NIOS NAT detection is designed to work with a specific form of NAT where multiple clients are NAT'ed onto a single public IP address using a single port block for each client at a time. Some NAT vendors refer to this scheme as PBA (Port Block Allocation) with fixed size port blocks. The NAT device cannot re-use the same NAT IP address for a different client for at least the drop interval of the threat protection rule that is currently running. Note that the client idle timeout duration of the NAT device must be larger than the drop interval of the threat protection rule that is currently running.

You can configure NAT Mapping rules on Advance Appliances by mapping a source IP address, network, or range of IP addresses with a range of ports and specifying the port block size to divide each port range into chunks of port blocks. Each port block represents a single NAT'ed client source port. The IP addresses and the port blocks specified in a NAT Mapping rule may be assigned to the clients in any order, either fixed, sequential, or random.

The appliance logs information about NAT'ed clients to the syslog. Following is an example of the threat detection event log message in the syslog for NAT'ed clients:

```
2015-06-01T22:57:22+00:00 daemon infoblox.localdomain threat-protect-  
log[12192]: err  
CEF:0|Infoblox|NIOS Threat|7.2.0-283371|120303001|Blacklist:block.com|7|  
src=3.0.0.100  
spt=1221 dst=1.1.6.2 dpt=53 act="DROP" cat="BLACKLIST UDP FQDN lookup" nat=1  
nfpt=1124  
nlpt=1223
```

The following values in the threat detection event log message are specific for NAT'ed clients:

- **nat=1**: Indicates that the syslog event is logged for a NAT'ed client.
- **nfpt**: Indicates the first port in the port block.
- **nlpt**: Indicates the last port in the port block.

For information about the syslog and how to use it, see [Viewing the Syslog](#). Note the following about the NAT Mapping feature:

- A single NAT'ed client cannot use multiple source IP addresses or multiple port blocks simultaneously, otherwise NIOS might consider the same NAT'ed client as different clients.
- NIOS does not support traffic originating from both NAT'ed and non-NAT'ed clients using the same source IP address.
- NAT mapping supports only IPv4 addresses, networks and ranges and any traffic originating from an IPv6 address, network, or range is considered as a non-NAT'ed client.
- NAT mapping is applicable only for traffic that uses UDP or TCP source port.



- NIOS cannot determine the pre-NAT IP address of a client. However, you can use the NIOS threat protection logs and the threat protection reports in conjunction with the NAT device logs for the same interval to determine the pre-NAT IP address of a NAT'ed client.
- The NAT mapping configuration in the NIOS appliance must be aligned with the actual NAT device configuration to ensure correct NAT detection.
- NAT mapping is not available for members that use Software ADP.

You can enable the NAT Mapping feature and configure NAT Mapping rules on the Infoblox-4030 10GE appliances and the following hardware-based appliances: PT-1405, PT-2205 and PT-2205-10GE. Note that you can enable this feature only when threat protection service is enabled on the appliance.

## Configuring NAT Mapping Rules

To enable the NAT Mapping feature and to configure NAT Mapping rules for a Grid or a member:

1. **Grid:** From the **Data Management** tab, select the **Security** tab, and then click **Grid Security Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **Security** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.
2. In the *Grid Security Properties* or *Member Security Properties* editor, select the **NAT Mappings** tab -> **Basic** tab, and then complete the following:  
The member inherits the NAT mapping settings from the Grid, by default. To override the settings, click **Override**. To retain the same settings as the Grid, click **Inherit**. Note that if you choose to override the Grid settings, you must override the entire setting; you cannot override partial setting.
  - **Enable tracking for NAT mappings:** This checkbox is disabled by default. When you select this checkbox, you can configure certain addresses, networks, and ranges for NAT tracking. When the NAT Mapping feature is enabled or disabled for the Grid, all members in the Grid inherit this setting.
  - **NAT Mappings:** To define the NAT Mapping rule, click the Add icon and select one of the following:
    - **Address:** Select this to configure a single NAT address. For example, if you have only one single NAT client or a single NAT source IP, select this.
    - **Network:** Select this if you want to configure an entire subnet for NAT.
    - **Range:** Select this if you want to configure a partial subnet for NAT.

After you select the NAT mapping type, the appliance adds a new row to the table and displays the type in the table (you cannot modify this in the table). In the **Address/Network/Range** column of the new row, enter a single NAT address, if you have selected **Address**. You can also configure a network address as a single NAT address. Enter a subnet if you have selected **Network**; or enter the partial subnet if you have selected **Range**. Note that IPv6 address, network, or range entry is not supported.

Click **+** beside the NAT mapping type, and then click the cell of the following fields in the new row to enter information:

- **Start Port:** Enter the first source port in the port range. For example, if the port range is 1000-2000, you can enter 1000 as the start port.
- **End Port:** Enter the end source port in the port range. For example, if the port range is 1000-2000, you can enter 2000 as the end port.
- **Block size:** Enter the port block size. This determines the number of port blocks and each port block is considered as a logical NAT client. Based on the port block size, the port range is divided into port blocks. For example, if the start port is 1000 and the end port is 2000, and if you enter the port block size as 100, then 1000-1099 is the first port block, 1100-1199 is the second port block, 1200-1299 is the third port block, and so on.

You can create a separate NAT Mapping rule for each source IP address, network, or range. You can add a maximum of 32 NAT Mapping rules and a maximum of 32 port ranges for each rule.



### Note

The combination of the source IP address, network, or range and the port blocks configured for each NAT'ed client must be unique and it should not overlap.

3. Save the configuration.

 **Note**

If you change the NAT Mapping rule after you have configured the NAT Mapping rule or if you enable or disable the NAT Mapping feature, you must publish the changes by clicking **Publish Changes** from the Toolbar. For information, see [Publishing Rule Updates](#).

## Managing Threat Protection Rulesets

You can do the following after the initial setup, including uploading the initial ruleset:

- Review the list of threat protection members, as described in [Listing Members](#).
- Look at rulesets that are currently installed on your system, as described in [Viewing Threat Protection Rules](#).
- Upload new rulesets to the system when you have selected to manually apply rule updates, as described in [Manually Uploading Rulesets](#).
- Publish rule updates that you have uploaded to the system, as described in [Publishing Rule Updates](#).
- Modify the **Comment** and **Do Not Delete** parameters for an existing ruleset, as described in [Modifying Rulesets](#).
- View differences between an old and a new rulesets and optionally merge rule parameter changes from the old ruleset into the new one, as described in [Comparing and Merging Rulesets](#).
- Configure proxy settings for the Grid if you need to use a proxy server for ruleset updates, as described in [Configuring Proxy Servers](#).
- Delegate ruleset updates to a Grid member and select an alternate interface for downloading the updates, as described in Configuring Members and Interfaces for Automatic Updates, see [Configuring Proxy Servers](#).

## Publishing Rule Updates

You can publish rule updates at any time after you have uploaded the ruleset. For information about uploading ruleset files, see [Manually Uploading Rulesets](#).

To publish rule updates:

1. From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab, and then click **Publish Changes** from the Toolbar.
2. In the *Publish Changes* dialog box, complete the following:
  - **Publish Changes on all Members:** Select one of the following:
    - **Simultaneously:** Publish changes on all of the members in the Grid at the same time.
    - **Sequentially:** Publish changes on each Grid member according to the number of seconds you enter in the **Sequential every (seconds)** field. For example, if you enter every 10 seconds, the system update changes on the first member, and 10 seconds later on the second member. This is the default option.
  - **Impacted Members and Services:** Click the Poll Members icon to display the affected members in that Grid. Grid Manager displays the member names and whether each member is configured for the threat protection service:
    - **Yes:** The service is active and the system will publish rule updates on this member upon execution of this task.
    - **No:** The service is not active and the system will not publish rule updates on this member.
    - **Disabled:** The service is currently disabled on this member.

To schedule this task, click the Schedule icon at the top of the dialog box. In the Schedule Change panel, complete the following:

- **Now:** Publish rule updates upon clicking **Publish**.
- **Later:** Enter the following information to schedule publishing updates at a certain date and time:
  - **Start Date:** Enter a date in YYYY-MM-DD (year-month-day) format. The appliance displays today's date. You can also click the calendar icon to select a date from the calendar widget.

- **Start Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. When you enter the time in a 24-hour format such as 23:00, Grid Manager displays 11:00:00 PM. You can also select a time from the drop-down list by clicking the time icon.
  - **Time Zone:** Select a time zone for the scheduled date and time from the drop-down list. This field displays the time zone of the browser that the admin uses to log in to Grid Manager.
3. Click **Publish** to publish changes immediately or click **Schedule Publish** to schedule the publish.

## Manually Uploading Rulesets

You can download a threat protection ruleset any time when you select to manually perform ruleset updates. You can choose to download rule updates but not immediately deploy them. NIOS archives and tracks up to nine ruleset versions, allowing for switching and merging between these versions when necessary. After uploading a ruleset, you can apply it by publishing it to the Grid and individual members. For more information, see [Publishing Rule Updates](#). Ruleset updates do not require restart of the DNS or threat protection service in the Grid, and they do not affect ongoing services. However, the appliance deploys updated rulesets only when you publish the changes. Note that all threat protection rule update events are logged in the syslog on the Grid Master only.

### Note

By default, threat protection ruleset updates are automatic. Infoblox recommends that you retain this setting. For information about how to configure this setting, see [Configuring Grid Security Properties](#).

To manually upload a ruleset file:

1. Access the Infoblox KB article # 2646 by logging in to the Infoblox Support site at <https://support.infoblox.com>, and then download the ruleset file in the KB article.
2. From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab, and then click the Add icon.
3. In the *Rule File Upload* dialog box, do the following:
  - **File:** Click **Select** to navigate to the ruleset file location, and then click **Upload**. Grid Manager displays the file name in this field.
  - Click **Test** to check the changes that will occur during the rule update, without actually applying the update. You can view the update details in the Syslog Viewer. The appliance preserves the uploaded file if you do not click **Update** to update the rules. When you manually upload rulesets the next time, this file will be displayed in the dialog. You can then choose to apply the update from this file or upload a new file before performing the update.
  - Click **Update** to update the rules.
  - Click **View Update Results** to view the updated rules in the Syslog Viewer. All threat protection rule updates are logged in the syslog on the Grid Master.

## Comparing and Merging Rulesets

After you manually download a ruleset and before you publish it, you can view differences between the old ruleset and the new one. The appliance shows you the system-level changes, including new rules, deleted rules, and rule syntax, between the two rulesets. It also shows you the customized parameter changes between the two versions. You can then select the changes you want to merge into the new ruleset. You can also modify some of these customized changes before you merge them into the new ruleset. Note that the modifications you make to the customized rule parameters will be added only to the new ruleset. When merging rulesets, all rules in the old ruleset and new ruleset are compared and identified by their rule IDs.

To merge rules from an older ruleset version to a newer ruleset:

1. From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab, click **Merge** from the Toolbar, and then select **Ruleset** from the drop-down list.
2. In the *Merge Changes Into Ruleset* editor, complete the following:
  - **Old Ruleset:** From the drop-down list, select the ruleset version from which you want to merge changes into the new ruleset. The **Comment** field displays additional information about the ruleset.

- **New Ruleset:** From the drop-down list, select the ruleset version to which you want the changes to be merged. The **Comment** field displays additional information about the ruleset.
- **Get Differences:** Click this to display a list of differences between the selected old and new ruleset versions. The appliance displays system-level changes in the **System changes from old ruleset (x items)** table, where x is the total number of changed rules between the old and new ruleset versions. The appliance merges all system-level changes listed in this table into the new ruleset.

The table displays the following information for each changed rule:

- **Rule ID:** The rule ID of the changed rule. You can click the rule link and modify parameters in the rule editor.
- **Rule Name:** The name of the rule.
- **Type:** Indicates whether the rule is a newly added rule or it has been deleted.
- **Old Ruleset Value:** Displays the old value that has been changed.
- **New Ruleset Value:** Displays the newly changed value.

The **Customizations from old ruleset (x items)** table displays customized rule parameter changes between the old and new rulesets, where x is the total number of changed rules. You can select all or specific changed rules in this table to be merged into the new ruleset. You can also modify parameter for selected rules before merging the changes into the new ruleset. This table displays the following information for each rule:

- **Member:** The Grid member on which this rule is currently running.
  - **Rule ID:** The rule ID of the changed rule. You can click the rule link and modify parameters in the rule editor.
  - **Rule Name:** The name of the rule.
  - **Old Ruleset Value:** Displays the old value that has been changed.
  - **New Ruleset Value:** Displays the newly changed value.
  - **Action:** Displays what the appliance will do to changes in this rule when you merge the rulesets.
3. Click **Merge Changes** to copy all the selected rules and changes you made to specific rules to the new ruleset. Note that you must select a rule from the **Customizations from old ruleset (x items)** table to activate the merging operation.

You can do the following in the *Merge Changes Into Ruleset* editor:

- Click **Export** to export all the changes listed in both tables to one CSV file. You can export this data after you click **Get Differences** and Grid Manager displays changes in both tables.
- Perform another ruleset merge by selecting another old ruleset and new ruleset.
- Click **Close** to exit the editor. Note that the editor does not close automatically after a merge.

## Modifying Rulesets

To modify parameters for an existing ruleset:

1. From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab.
2. In the Grid Threat Ruleset table, select the ruleset checkbox, and then click the Edit icon.
3. In the *Threat Ruleset* editor, select the **General** tab -> **Basic** tab to modify the following:
  - **Comment:** Enter information about the ruleset.
  - **Mark as Do Not Delete:** When you select this, the ruleset cannot be deleted during an update. You can select up to three rulesets and mark them as **Do Not Delete**.

You can also view the following information (but you cannot modify it):

- **Version:** Displays the ruleset version in YYYYMMDD-x format, where YYYY is the year, MM the month, DD the date, and x the ruleset engine version number.
- **Active On:** Displays the Grid or the name of the Grid member on which the ruleset is currently running. This can be **Grid**, a Grid member name, or **None**. When this displays **None**, the ruleset is not being used.
- **Added On:** Displays the timestamp when the ruleset was uploaded to the Grid Master in this format: YYYY-MM-DD HH:MM:SS, plus time zone.

## Overriding the Grid Ruleset

You can override the Grid ruleset for individual members. Note that when a member uses a different ruleset than the Grid, it does not receive a ruleset update if the Grid automatically receives the next update. For information about how to activate a ruleset for the Grid, see [Configuring Grid Security Properties](#).


To override the Grid ruleset:

1. From the **Data Management** tab, select the **Security** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.
2. In the *Member Security Properties* editor, select the **Threat Protection** tab -> **Basic** tab, and complete the following:
  - **Use ruleset:** Click this radio button to edit the ruleset. You can inherit or override a ruleset. When you select to use a ruleset, select an active ruleset from the **Threat Protection Ruleset Version** drop-down list.
  - **Use profile:** Click this radio button to edit the threat protection profile. Click **Select Profile** to select a threat protection profile from the *Threat Protection Profile Selector* dialog box. For more information about threat protection profiles, see [Adding Threat Protection Profiles](#).
  - **Active Ruleset Version:** Click **Override** and select a ruleset from the drop-down list. The appliance activates and applies this ruleset to the member. In the confirmation dialog, click **Yes** and the member switches to the selected ruleset. This field is read-only at the member level.
  - **Active Ruleset Comment:** Displays information about the selected ruleset from the **Comment** field. This field is read-only at the member level.

## Activating a Ruleset

The appliance can store up to nine (9) versions of rulesets. You can switch between these versions and activate one of them as the default ruleset.

To activate a specific ruleset:

1. **Grid:** From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab.  
**Member:** From the **Data Management** tab, select the **Security** tab -> **Members** tab -> *member* link.
2. In the Grid Threat Rulesets or Member Rulesets table, click the Action icon  and select **Activate** from the menu. In the confirmation dialog, select **Yes** to proceed. The appliance installs the ruleset and uses it as the default ruleset for the Grid or the member.

## Viewing Threat Protection Rulesets

Grid Manager displays all the rulesets that have been automatically or manually uploaded to the Grid Master. You can drill down to each ruleset to review individual rules in each rule category. To view the current threat protection rulesets and rules:

1. **Grid:** From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab.  
**Member:** From the **Data Management** tab, select the **Security** tab -> **Members** tab -> *member* link.  
**Profile:** From the **Data Management** tab, select the **Security** tab -> **Profiles** tab -> *profiles* link.
2. In the **Threat Protection Rules** or the **Members** tab or the **Profiles** tab, Grid Manager displays current rulesets that are stored in the system. It displays the following information for each ruleset:




### Note

If the member is associated with a profile, the rules become read-only and you cannot edit them when you drill down the ruleset from the **Members** tab. However, you can still edit the rules if you drill down the ruleset from the **Profiles** tab even though Grid members are associated with the selected profile.

- **Version:** Displays the ruleset version in YYYYMMDD-x format, where YYYY is the year, MM the month, DD the date, and x the rule engine version number. You can click a version link to view rule categories and individual rules in each category for that specific ruleset.

- **Active On:** Displays the Grid or the name of the Grid member on which the ruleset is currently running. This can be **Grid**, a Grid member name, or **None**. When this displays **None**, the ruleset is not being used.
- **Do Not Delete:** Indicates whether this ruleset can be deleted or not during a ruleset update.
- **Added On:** Displays the timestamp when the ruleset was uploaded to the Grid Master in this format: YYYY-MM-DD HH:MM:SS, plus time zone.
- **Add Type:** Indicates whether the ruleset was uploaded automatically or manually.
- **Comment:** Additional information about this ruleset.

You can also do the following when you left click the Action icon :

- Mark a ruleset as do not delete.
- Activate, open, or edit a specific ruleset.
- Click **Reset to defaults** to globally reset all the rule definitions to their default settings in the selected ruleset. When you reset the ruleset to defaults, all member rulesets will be reset as well. In the *Ruleset Reset* dialog, you can also select **Delete all custom rules in selected ruleset** to remove all the custom rules in the ruleset. You can reset to defaults at the rule category and rule levels. Note that the **Reset to defaults** option is not available if you upgrade to NIOS 7.3.x from a previous release.

To view rules in a specific ruleset version, click the **Version** link and Grid Manager displays the threat protection rules by categories. You can also select the checkbox of a ruleset and click the Open icon to view the rules in the ruleset.

You can also do the following in this panel:

- Manually upload a ruleset by clicking the Add icon. For more information, see [Manually Uploading Rulesets](#).
- Modify some of the data in the table. Double click a row, and you can modify the **Do Not Delete** and **Comment** columns. Click **Save** to save the changes. Note that other fields are read-only.
- Select the checkbox of a ruleset and click the Delete icon to delete a ruleset, if it is not marked as **Do Not Delete**.
- Print or export the data.




#### Note

When you delete a ruleset that was recently downloaded through automatic downloads, the appliance cannot retrieve or automatically download this ruleset version again. You must manually download the ruleset and then manually deploy it to the Grid. For information about manual downloads, see [Manually Uploading Rulesets](#).

## Listing Members

This panel displays the following information about each Grid member:

- **Name:** The hostname of the Grid member.
- **Status:** The status of the DDoS service running on the Grid member.
- **Version:** Displays the threat protection ruleset version that is currently running on this member. This is displayed in YYYYMMDD-x format, where YYYY is the year, MM the month, DD the date, and x the rule engine version number.
- **Comment:** Comments that were entered for the Grid member.
- **Profile:** Displays the threat protection profile that is associated with this member.
- **Site:** Values that were entered for this pre-defined attribute. You can do the following:
  - Sort the data in ascending or descending order by column.
  - Select a member and click the Edit icon to modify the data.
- Click the Action icon  next to the respective member and:
  - select **Permissions** from the menu to define permissions for the member.
  - select **Clone to new profile** from the menu to create a new threat protection profile by cloning the settings from a member to the new profile. Note that this option is enabled only when the selected member is not associated with a profile.
- Print and export the data in this tab.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria:
  - a. In the filter section, click **Show Filter** and define filter criteria for the quick filter.
  - b. Click **Save** and complete the configuration In the Save Quick Filter dialog box.

The appliance adds the quick filter to the quick filter drop-down list in the panel. Note that global filters are prefixed with [G], local filters with [L], and system filters with [S].

## Delegating Rule Updates to Grid Members

You can delegate ruleset updates to other Grid members if desired. You can also configure the interface you want to use for automatic updates on the members. For information about how to configure the Members and Interfaces for Automatic Updates, see [Configuring Proxy Servers](#).

## Configuring Threat Protection Profiles

When you configure the Grid or Member security properties, you have an option to select an active ruleset or a threat protection profile. A threat protection profile defines specific security settings and ruleset that you want to apply to a specific member or a group of members. Infoblox uses event filters to limit the amount of logs that are generated due to the threat protection events. It drops packets for services or ports that are not enabled on the card.

If you want to use the same threat protection rulesets and settings for multiple members, you can create a threat protection profile and associate it with multiple members so you do not have to configure them individually. You can configure threat protection profiles for both hardware and Software ADP appliances.

Infoblox automatically migrates threat protection profile settings when:

- you update an active ruleset at the Grid level and if a profile has inherited the respective active ruleset from the Grid.
- the profile inherits an active ruleset from the Grid and you override the active ruleset that is associated with the profile.
- the profile is set to override an active ruleset and you change an active ruleset for the profile.

You can do the following to configure threat protection profiles:

- Create threat protection profiles by associating rulesets, event filters and configuring multiple DNS requests over a single TCP session, as described in [Adding Threat Protection Profiles](#).
- Clone threat protection profiles, as described in [Cloning Threat Protection Profiles](#).
- Modify threat protection profiles, as described in [Modifying Threat Protection Profiles](#).
- View differences between an old and a new rulesets and merge changes from an old threat protection profile into the new one, as described in [Merging Threat Protection Profiles](#).
- Inherit Grid rule settings for a threat protection profile, as described in [Inheriting Grid Rule Settings](#).
- Delete a threat protection profile, as described in [Deleting Threat Protection Profiles](#).
- View the list of threat protection profiles, as described in [Viewing Threat Protection Rulesets](#).

## Viewing Threat Protection Profiles

To view the list of threat protection profiles:

1. From the **Data Management** tab, select the **Security** tab, and then click the **Profiles** tab. Grid Manager displays the following information:
  - **Name:** The name of the threat protection profile. Click the *profile* name to view the rules that are associated with the selected threat protection profile. For more information about rules, see [Viewing Threat Protection Rulesets](#).
  - **Version:** The version of the threat protection profile.
  - **Members:** Displays the member associated with the threat protection profile. You can sort the values in this column.
  - **Comment:** Displays information that you specified for the threat protection profile.
  - **Site:** The location to which the member belongs. This is one of the predefined extensible attributes.

The **Disabled** column indicates whether the rule is enabled or disabled. The **Rule ID** column is set to grey for all rules that are disabled. The blue background for some parameters indicates that these parameters have been inherited from the Grid ruleset.



You can also do the following:


- Use **Global Search** to search for threat protection profiles by name, comment or object type. For information, see [Using Global Search](#).
- Use **Smart Folders** to organize threat protection profiles by name, comment or object type. For information, see [Smart Folders](#).
- To export the entire list of threat protection profiles in a csv format, click the Export icon and choose **Export Data in Infoblox CSV Import Format**. It also exports the customized rules for the profile. For more information, see [Exporting Data to Files](#). To export all data in a different format, click the Export icon and choose **Export Visible Data**. For more information, see [Exporting Displayed Data](#).
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria:
  - a. In the filter section, click **Show Filter** and define filter criteria for the quick filter.
  - b. Click **Save** and complete the configuration in the *Save Quick Filter* dialog box.

The appliance adds the quick filter to the quick filter drop-down list in the panel. Note that global filters are prefixed with [G], local filters with [L], and system filters with [S].


## Deleting Threat Protection Profiles

When you select this option, NIOS deletes the selected threat protection profile from the database. When you delete a threat protection profile or a profile rule, the appliance puts them in the Recycle Bin, if enabled. You can restore them if necessary. Note the following about threat protection profiles:

- When you delete a profile that contains an active ruleset, which is overridden at the profile level, you can restore the profile including the configuration if the ruleset exists on the Grid. You cannot restore the profile if the ruleset is deleted.
- When you delete a profile that inherits ruleset from the Grid, and if the same ruleset is still active on the Grid, you can restore the profile including the configuration for profile rules. If the ruleset is deleted or if the active ruleset on the Grid has changed, the appliance restores the profile from the Grid along with the profile rule settings for the current active ruleset.

To delete a threat protection profile, from the **Data Management** tab, select the **Security** tab -> **Profiles** tab, select the threat protection profile that you want to delete, click the Action icon  and then select **Delete**. Click **Yes** in the confirmation dialog box to delete.

## Inheriting Grid Rule Settings

When you select this option, NIOS automatically inherits the rule settings of the Grid for the respective profile. For example, if you update a rule at the Grid level, you can select this option to reflect the same changes in the profile. To inherit Grid rule settings for the selected profile, from the **Data Management** tab, select the **Security** tab -> **Profiles** tab, select the threat protection profile, click the Action icon  and then select **Inherit Grid Rule Settings**. Click **Yes** in the confirmation dialog box to inherit the settings.

## Merging Threat Protection Profiles

You can merge the threat protection rulesets from one profile to another before you publish it. The following rules are applicable when you update a ruleset that is assigned to a profile and migrate the profile settings to a new ruleset:

- NIOS migrates the profile settings of the rule to the new ruleset when you override a rule, which exists in both the old and the new ruleset, at the profile level.
- The profile ruleset continues to inherit the settings from the Grid if you do not override a rule that exists in both the old and the new ruleset.
- NIOS removes the profile settings for a rule when you override the rule settings for a profile and the rule does not exist in the new ruleset.
- When a rule exists only in a new ruleset, it inherits the settings from the Grid by default.
- Parameters for custom rules are migrated only when the same custom rule exists in the new ruleset.



You can view differences between the old profile and the new one before you publish the changes. The appliance shows you the system-level changes, including new rules, deleted rules, and rule syntax, between the two rulesets and customizations from the old profile. You can then select the changes you want to merge into the new profile. You can also modify some of these customized changes before you merge them into the new profile. When merging profiles, all rules in the old ruleset and new ruleset are compared and identified by their rule IDs.

To merge a threat protection profile:

1. From the **Data Management** tab, select the **Security** tab -> select either **Threat Protection Rules** or **Members** or **Profiles** tab, click **Merge** from the Toolbar, and then select **Profile** from the drop-down list.
2. In the *Merge Changes Into Profile Ruleset* editor, complete the following:
  - **Old Profile:** From the drop-down list, select the threat protection profile from which you want to merge changes into the new profile. The **Comment** field displays additional information about the profile and **Ruleset version** displays the ruleset version.
  - **New Profile:** From the drop-down list, select the threat protection profile to which you want to merge changes. The **Comment** field displays additional information about the profile and **Ruleset version** displays the ruleset version.
  - **Get Differences:** Click this to display a list of differences between the old and new profiles. The appliance displays system-level changes in the **System changes from old profile (x items)** table, where x is the total number of changed rules between the old and new profiles. The appliance merges all system-level changes listed in this table into the new profile.

The table displays the following information for each changed rule:

- **Rule ID:** The rule ID of the changed rule. You can click the rule link and modify parameters in the rule editor.
- **Rule Name:** The name of the rule.
- **Type:** Indicates whether the rule is a newly added rule or it has been deleted.
- **Old Ruleset Value:** Displays the old ruleset value.
- **New Ruleset Value:** Displays the new ruleset value.

The **Customizations from old profile (x items)** table displays customized rule parameter changes between the old and new profile, where x is the total number of changed rules. You can select all or specific changed rules in this table to be merged into the new profile. You can also modify the parameters for selected rules before merging the changes into the new profile. This table displays the following information for each rule:

- **Member/Profile:** The Grid member on which this rule is currently running.
  - **Rule ID:** The rule ID of the changed rule. You can click the rule link and modify parameters in the rule editor.
  - **Rule Name:** The name of the rule.
  - **Old Ruleset Value:** Displays the old ruleset value.
  - **New Ruleset Value:** Displays the new ruleset value.
  - **Action:** Displays what the appliance will do to changes in this rule when you merge the profiles.
3. Click **Merge Changes** to copy all the selected rules and changes you made to specific rules to the new profile. Note that you must select a rule from the **Customizations from old profile (x items)** table to activate the merging operation.


You can do the following in the *Merge Changes Into Profile Ruleset* editor:

- Click **Export** to export all the changes listed in both tables to one CSV file. You can export this data after you click **Get Differences** and Grid Manager displays changes in both tables.
- Perform another profile merge by selecting another old profile and new profile.
- Click **Close** to exit the editor. Note that the editor does not close automatically after a merge.

## Modifying Threat Protection Profiles

You can modify the details associated with a threat protection profile. You can edit the rules when you drill down to the ruleset even though a member is associated with the profile.

To modify a threat protection profile:

1. From the **Data Management** tab, select the **Security** tab -> **Profiles** tab, select the threat protection profile that you want to modify, click the Action icon  and then select **Edit**.

2. The *Threat Protection Profile* editor contains the following tabs from which you can modify information:
  - **General:** All fields are automatically propagated with available information. You can modify the values in the **Basic** and **Advanced** tabs. For more information, see [Adding Threat Protection Profiles](#).
  - **Member Assignment:** Add or delete members that are associated with the respective threat protection profile. Click the Add icon to associate a member with the selected profile. In the *Threat Protection Member Selector* dialog box, select the member you want to associate with the profile. For information, see [Listing Members](#). To delete a member that is associated with the profile, select the checkbox next to the respective member and click the Delete icon.
  - **Extensible Attributes:** Add and delete extensible attributes that are associated with the template. You can also modify the values of the extensible attributes. For information, see as described in [Modifying Extensible Attributes](#).
3. Save the configuration.

You must publish the changes after modifying a threat protection profile. For more information about publishing changes, see [Publishing Rule Updates](#).


 **Note**

By default, all ruleset versions, events per second per rule, and disable multiple DNS requests via single TCP session are inherited from the Grid unless you click **Override** to change the Grid settings.

## Cloning Threat Protection Profiles

You can create a new threat protection profile by cloning an existing one. The appliance creates a new profile with the settings that are copied from the source profile. You can also associate extensible attributes with the profile. Note that members in the source profile are not carried over to the cloned profile. You must associate new members with the newly cloned profile.

To clone a threat protection profile:

1. From the **Data Management** tab, select the **Security** tab -> **Profiles** tab, select the threat protection profile that you want to clone, click the Action icon  and select **Clone**.
2. In the *Clone Threat Protection Profile Wizard*, enter a name for the threat protection profile and modify the details. For more information, see [Adding Threat Protection Profiles](#).
3. Associate extensible attributes and members with the threat protection profile.
4. Save the configuration.

## Adding Threat Protection Profiles

You can create a threat protection profile and associate an active ruleset with it. Infoblox supports common threat protection rulesets for both hardware and Software ADP members. You can either upload a ruleset or download rulesets from a server. You can create any number of threat protection profiles, but you can select only a maximum of five rulesets in combination at the Grid, member and profile levels. For more information about rulesets, see [Understanding Threat Protection Rulesets and Rules](#).

The threat protection profile allows you to create your own set of rules for either a member or a group of members that experience a similar kind of traffic. After you define a profile, you can clone it and test the copied settings for a new ruleset on one member before publishing the changes for a group of members that are associated with the profile.

To define threat protection profiles:

1. From the **Data Management** tab, select the **Security** tab -> **Profiles** tab and then click the *Add* icon.
2. In the *Add Threat Protection Profile Wizard*, add the following:
  - **Name:** Enter a name for the threat protection profile.
  - **Comment:** Enter information about the threat protection profile.
  - **Active Ruleset Version:** Select a value from the drop-down list. This indicates the current ruleset that is used for the respective threat protection profile. If you inherit a ruleset from the Grid and later change the respective ruleset at the Grid level, the new ruleset is not reflected in the profile. You must manually change the selected ruleset for the profile. For more information about active rulesets, see [Understanding Threat Protection Rulesets and Rules](#).

- **Active Ruleset Comment:** Click **Override** to override the comment.
  - **Events per Second per Rule:** Click **Override** to override the values. This indicates the number of events that is logged per second per rule to allow the appliance to log events to the syslog. Specify the number of events logged per second per rule. The default value is one and the maximum value is 700. Setting the value to 0 (zero) disables the appliance from logging events for the rules. The appliance displays an error message when you enter a value greater than the maximum value. You can override this event filter at the member level. For more information, and guidelines about using this setting, see [Using the Events Per Second Rule Setting](#).
  - **Disable multiple DNS requests via single TCP session:** Click **Override** to override the values. This determines if multiple DNS responses through TCP connection are disabled. For more information, see [Enabling Multiple DNS Requests through a Single TCP Session](#).
3. Click **Next** to add extensible attributes.
  4. Save the configuration.



#### Note

A member associated with a threat protection profile can neither modify **Events per Second per Rule** and **Disable multiple DNS requests via single TCP session** settings at the member level nor enable or disable rules and change rule parameters at the member level.

## Managing Threat Protection Rules

You can modify any previously defined custom rules, or some of the parameters for system and auto rules. For most system and auto rules, you may change the **Action** and **Log Severity**. You can also enable or disable individual rules or an entire category of rules.

If you have selected to manually update threat protection rules, you must download updated rules from the Infoblox Support web site and then publish them to the system.

When a member is associated with a profile, it automatically uses the ruleset that is associated with the profile. When you delete the associated profile, member uses the ruleset that was previously associated with it.

You can do the following after the initial setup, including uploading the initial ruleset:

- Look at rules that are currently installed on your system, as described in [Viewing Threat Protection Rulesets](#).
- Enable and disable certain rules, as described in [Enabling and Disabling Rules](#).
- Upload rule updates to the system when you have selected to manually apply rule updates, as described in [Modifying System and Auto Rules](#).
- Publish rule updates that you have uploaded to the system, as described in [Publishing Rule Updates](#).
- Modify existing system rules, as described in [Modifying System and Auto Rules](#).
- Modify custom rules, as described in [Modifying Custom Rules](#).

## Modifying Custom Rules

1. From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab.
2. In the Grid Rules table, expand the category to which the rule belongs, select the checkbox, and then click the Edit icon.
3. In the Custom Rule editor, select the **General** tab -> **Basic** tab to modify the following:
  - **Comment:** Enter information about the custom rule.
  - **Disable:** Select this checkbox to disable the custom rule.

You can also view the following information (but you cannot modify it):

- **Template:** Displays the name of the template the custom rule uses.
  - **Rule ID:** Displays the rule ID of the custom rule.
  - **Name:** Displays the name of the rule.
  - **Category:** Displays the category to which the custom rule belongs.
  - **Description:** Displays the description of the custom rule.
4. In the Custom Rule editor, select the **Settings** tab -> **Basic** tab to modify the following:

- **Log Severity:** Select the log severity level from the drop-down list. You can select **Critical, Major, Warning,** or **Informational**. Log severity may have an effect on how other Grid services respond to particular events. The selection here corresponds to the severity levels you can configure for logging in the syslog.
  - **Rule Parameters:** In the Rule Parameters table, the **Description** column displays the rule parameters. Click the row and enter the corresponding values for the rule parameters in the **Value** column.
  - **Action:** Displays the operation which the appliance performs when this event occurs. Some rules are restricted to specific actions. For example, the action for all blacklisting rules is set as **Drop** where the appliance drops the packets and logs the activity when such an event occurs. The action for all whitelisting rules is set as **Pass**, where the appliance silently passes the packets without logging when such an event occurs.
5. Save the configuration.


## Modifying System and Auto Rules

1. From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab.
2. In the Grid Rules table, expand the category to which the rule belongs, select the checkbox, and then click the Edit icon.
3. In the *System Rule* or *Auto Generated Rule* editor, select the **General** tab -> **Basic** tab to modify the following:
  - **Comment:** Enter information about the system rule.
  - **Disable:** Select this checkbox to disable the system rule. You cannot disable auto rules. You can also view the following information (but you cannot modify it):
  - **Rule ID:** Displays the rule ID of the system rule.
  - **Name:** Displays the name of the rule.
  - **Category:** Displays the category to which the system rule belongs.
  - **Description:** Displays the description of the system rule.
  - **Order:** Displays the number that indicates the order in which the rule will be executed by the appliance. The rule order can change during a ruleset update.
4. In the *System Rule* or *Auto Generated Rule* editor, select the **Settings** tab -> **Basic** tab. Depending on the rule, you may or may not be able to modify the following:
  - **Action:** Displays one of the following: **Alert, Drop** or **Pass**. Some rules are restricted to specific actions. For example, the action for all blacklisting rules is set as **Drop**, where the appliance drops the packets and logs the activity when such an event occurs. The action for all whitelisting rules is set as **Pass**, where the appliance silently passes the packets without logging when such an event occurs.
    - **Alert:** Logs the activity, and passes the packet.
    - **Drop:** Logs the activity and drops the packet.
    - **Pass:** Silently passes the packet without logging.
  - **Log Severity:** Select the log severity level from the drop-down list. You can select **Critical, Major, Warning,** or **Informational**. The selection here corresponds to the severity levels you configure for logging in the syslog.
  - **Rule Parameters:** In the Rule Parameters table, the **Description** column displays the rule parameters. Click the row and enter the corresponding values for the rule parameters in the **Value** column. Depending on the rule, this table displays only the parameters that are relevant to the system or auto rule.
5. Save the configuration.

## Enabling and Disabling Rules

By default, all activated threat protection rules apply across the entire Grid. Enabling or disabling a rule category will enable or disable all rules contained in that category. You can also enable or disable individual rules.

To enable or disable all rules in a category, do the following:

1. Grid: From the **Data Management** tab, select the **Security** tab -> **Threat Protection Rules** tab, and then click the ruleset link.  
Member: From the **Data Management** tab, select the **Security** tab -> **Members** tab, and then select a member and click the ruleset link.
2. In the Threat Ruleset table, click the Action icon  next to the rule category, and choose **Enable All Rules in Category** or **Disable All Rules in Category**. Either option can be disabled depending on the current state of the rules in the category. For members, you can select **Inherit Grid Rule Settings in Category** to inherit rule settings from the Grid.

To enable or disable individual rules, do the following:

1. Grid: From the **Data Management** tab, select the **Security** tab → **Threat Protection Rules** tab, and then click the ruleset link.  
Member: From the **Data Management** tab, select the **Security** tab → **Members** tab, and then select a member and click the ruleset link.
2. Click the arrow beside a rule category to expand all rules in a rule category.
3. Click the Action icon next to a rule and choose **Enable** or **Disable** from the menu. Either option can be disabled depending on the current state of the rule.



#### Note

Depending on the nature of the rules, you may or may not be able to disable or enable certain rules.

## Viewing Threat Protection Rules

You can view the threat protection rules in one of the following views:

- Click **Toggle Flat View** to display a flat list of all the threat protection rules. In the flat view, each of the custom, system, and auto rules are listed separately.
- Click **Toggle Tree View** to display only the category of threat protection rules. You can expand the category of rule to view individual rules listed in each category.

To view rules categories and individual rules in a specific ruleset version:

1. **Grid:** From the **Data Management** tab, select the **Security** tab → **Threat Protection Rules** tab. **Member:** From the **Data Management** tab, select the **Security** tab → **Members** tab → *member* link. **Profile:** From the **Data Management** tab, select the **Security** tab → **Profiles** tab → *profile* link.
2. In the Grid Security or Member Security table, click the **Version** link and Grid Manager displays the threat protection rules by categories. The **Category** column lists all the category to which rules belong.
3. To view individual rules listed in each category, expand the list by clicking the arrow beside the checkbox. You can view the following information for each rule:
  - **Category:** The category to which the rule belongs.
  - **Order:** The number that indicates the order in which the rule will be executed by the appliance. The rule order can change during a ruleset update.
  - **Rule ID:** The ID of the rule.
  - **Rule Name:** The name of the rule. This can contain up to 255 characters.
  - **Type:** The rule type. This can be **Custom**, **System**, or **Auto**. For more information about each rule type, see [Understanding Threat Protection Rulesets and Rules](#).
  - **Disabled:** Displays whether the rule is disabled. A disabled rule does not perform any mitigation functions.
  - **Comment:** Comments that were entered for the rule. This can contain up to 255 characters.
  - **Action:** The operation that the appliance performs when the event occurs. This can be one of the following:
    - **Alert:** The appliance passes the packets and logs the event.
    - **Drop:** The appliance drops the packets and logs the event.
    - **Pass:** The appliance passes the packets but does not log the event.
  - **Description:** Description about the rule. This can contain up to 255 characters.
  - **Rule Parameters:** Displays the rule parameters that are configured for the rules and the corresponding values for the rule parameters.
  - **Log Severity:** Log severity level. This can be **Critical**, **Major**, **Warning**, or **Informational**. You can also do the following in this panel:
    - Click the Action icon
    -

and select one of the following actions for a rule category:

- **Enable All Rules in Category:** Select this to enable all the rules in the selected category. For a Grid member, this action overrides the Grid rule settings.
- **Disable All Rules in Category:** Select this to disable all rules in the selected category. For a Grid member, this action overrides the Grid rule settings.

- **Inherit Grid Rule Settings in Category:** Select this to inherit Grid rule settings for the selected category. This appears only for member settings.
- Modify some of the data in the rules table. Double click a row, and modify the data. Click **Save** to save the changes. Note that some fields are read-only.
- Select the checkbox of a rule and click the Edit icon to modify the properties of the rule.
- Select the checkbox of a custom rule and click the Delete icon to delete a custom rule.
- Print or export the data.
- Publish changes you make to the rules by clicking **Publish Changes** from the Toolbar. For more information, see [Publishing Rule Updates](#).
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.



#### Note

When you use the **Go to** function to search for specific rules, Grid Manager may return duplicates of the same rules due to the paging mechanism currently used for displaying information.

- Create a quick filter to save frequently used filter criteria. For more information about using quick filters, see the following section.

### Using Quick Filters

You can use the following quick filters to filter threat protection rules in the Grid Rules panel. When you select a filter, Grid Manager displays only the specified rules. Using filters makes it easier to locate specific rules for editing, enabling, and disabling.

- **All Auto Generated Rules:** This option shows all auto-generated rules defined in the Infoblox Advanced DNS Protection solution.
- **All Custom Rules:** This option shows all whitelisting and blacklisting custom rules defined by NIOS users.
- **All System Rules:** All protocol-specific rules associated with particular attack phenomena.
- **All Disabled Rules:** This option shows all disabled rules.
- **All Enabled Rules:** This option shows all enabled rules.

For more information about filters, see [Finding and Restoring Data](#).

## Monitoring Threat Protection Events

You can monitor threat protection events through the following:

- Syslog, as described in [Monitoring through Syslog](#).
- Status Dashboard, as described in [Threat Protection Statistics Widget](#).
- Threat protection report, as described in [Threat Protection Reports](#).

### Threat Protection Reports

NIOS provides a series of reports to monitor and analyze DNS threat protection events. When you integrate a reporting member in your Grid, you can get the threat protection both hardware and Software ADP related reports so you can monitor event counts by severity, member, rule, and rule category. Note that all reporting traffic goes over the MGMT port on the reporting member by default. If for any reasons you need to configure a different interface for the reporting traffic, you can do so through the reporting *Member Properties* editor. For detailed information about these reports and how to configure the interface for reporting traffic, see [Security Dashboards](#).

## Threat Protection Statistics Widget

You can also get a high-level view of the threat protection events through the Status Dashboard. The Advance Appliances provide the *Threat Protection Statistics* widget so you can monitor the trend and counts of the various events. For more information about the Dashboard and the *Threat Protection Statistics* widget, see [Status Dashboard](#).

## Monitoring through Syslog

To receive threat protection events in the syslog, you must enable the Security option in the DNS logging category of the Grid DNS Properties editor. For information about configuring the logging category as described in [Setting DNS Logging Categories](#), see [Using a Syslog Server](#). Once the Security option is enabled, hardware-based appliances log each threat protection related event in the syslog in CEF (Common Even Format). You can get detailed information about the events by reviewing the syslog periodically. For information about how to configure the syslog server, see [Using a Syslog Server](#).

When a DNS attack is detected against an enabled rule, the appliance generates a log message. Note that only threat protection messages in CEF are displayed in the syslog. The log messages for rate limiting alert events also include the FQDNs extracted from DNS queries whose standard query and question count is greater than zero so you can quickly identify the offending clients. Note that the FQDN field displays "NA" for invalid DNS queries. This feature is enabled by default. You can disable this only in Maintenance Mode using the CLI command `set smartnic-debug-adp-log-fqdn off`.

Example:

When the appliance detects ICMP ping attacks that exceed the pint size against an existing auto rule that has the following configuration:

```
Log Severity = Critical
Rule ID = 120600925
Rule Name = Potential DDoS related domain
Rule Action = Drop
Rule Category = Potential DDoS related Domains
```

It generates the following threat detection event log message:

```
2018-04-20T09:43:21+00:00 daemon infoblox . localdomain named[14792]: info CEF:0|
Infoblox|NIOs|8.3.0-369415|RPZ-QNAME|Local-Data|7|app=DNS dst=10.34.173.11
src=10.120.20.28 spt=52240 view=_default qtype=A msg="rpz QNAME Local-Data
rewrite a_rec [A] via a_rec.local.com" IPSD=N/A Acct-Session-
Id=8333332d-11111111 Parental-Control-Policy=010000000033 Calling-Station-
Id=1101202041 NAS-PORT=1813 Subscriber-Secure-Policy=00000fff Guest=1
LocalID=000C2987FEEE CAT=RPZ
```

The number of log messages generated is based upon your Event per Second per Rule setting. For example, if the setting is 5, the appliance generates five log messages of the same event per second when the attack continues within the time duration. Each log message contains the following information:

- The timestamp when the event happened in yyyy-mm-ddThh:mm:ss+00:00 format.
- **Infoblox|NIOs|x.x.x**: Indicates the Infoblox product, and x.x.x represents the NIOS version.
- The string following the NIOS version is a hard-coded constant. In this example, it is RPZ QNAME.
- The number following the rule ID is the log severity. The following numbers indicate the severity levels:
  - **8 = Critical**
  - **7 = Major**
  - **6 = Warning**
  - **4 = Informational**
- **app**: DNS.



- **dst**: Destination IP address.
- **src**: Source IP address.
- **spt**: Source port.
- **view**: DNS view.
- **qtype**: Query type.
- **msg**: RPZ rule.
- **IPSD**: IP space discriminator.
- **Acct-Session-Id**: Session ID.
- **Parental-Control-Policy**: Parental Control Policy.
- **Calling-Station-Id**: Subscriber ID.
- **NAS-PORT**: NAS Port.
- **Subscriber-Secure-Policy**: Subscriber Secure Policy.
- **Guest**: Guest indicator. For fixed line or home router deployments, a guest indicator value '1' indicates guest device and '0' indicates subscriber device.
- **LocalID**: MAC address of the subscriber device. For fixed line or home router deployments, if the guest indicator value in the **Guest** field displays '1' then the **Local ID** field displays the MAC address of the guest device.
- **CAT**: The category to which the rule belongs. In this example, the category is "RPZ."

To view DNS threat protection related log messages:

1. From the **Administration** tab, select the **Logs** tab -> **Syslog** tab.
2. From the drop-down list at the upper right corner, select the Grid member on which you want to view the syslog.
3. From the Quick Filter drop-down list, select **Threat Rule Update Events** or **Threat Detection Event Logs** to view rule update events or threat detection events respectively. To narrow down the system messages you want to view, click Show Filter and then select the filters you want to use. For information about how to use filters, see [Using Filters](#).

## DNS and Network-Flood Threats

DNS is a tempting target for attacks given that it is a core Internet service. Attackers can send malformed DNS queries or DNS responses to the targeted server, hoping to exploit bugs in its DNS implementation. Other variants include code insertion, buffer overflows, memory corruption, NULL pointer dereferencing, and specific vulnerability exploits. DNS attacks tend to follow specific patterns but can be difficult to deal with using only rate-limiting techniques, because of the sheer scale of many recent attacks. DNS threat protection is designed to grow and expand over time, through threat protection rule updates, to deal with both outside-in and inside-out attacks on network infrastructure and Internet services.

Following are some of the network-flood attacks that can target your DNS caching and authoritative servers:

- [Internet Control Message Protocol \(ICMP\) Flood](#)
- [SYN Flood](#)
- [UDP DNS Flood](#)
- [Inside-Out Attacks](#)
- [DNS Cache Poisoning](#)
- [DNS Reflection and Amplification Attacks](#)
- [DNS Malware](#)
- [DNS Domain Hijacking](#)

### DNS Domain Hijacking

Domain hijacking or domain theft is the act of changing the registration of a domain name without the permission of its original registrant. Domain hijacking is generally done by exploiting a vulnerability in the domain name registration system or through social engineering. In some cases however, domain hijackers alter the DNS data of a domain after gaining control of it. They consequently redirect users to a fraudulent site, instead of the legitimate site, on the Internet. To protect your authoritative DNS server against DNS domain hijacking, you can configure the appliance to monitor NS records and glue records of top-level authoritative zones. Based on your configuration, the appliance periodically checks DNS data in the NS RRsets for these zones and compares the data with that in the appliance database. It then reports any data discrepancies through SNMP traps and logs related events in the syslog. You can also monitor the status of DNS data discrepancies, if any, through the DNS integrity check widget, see [Status Dashboard](#). The severity in data



discrepancies can help identify possible domain hijacking. For more information about how to configure this feature, see [Configuring DNS Integrity Check for Authoritative Zones](#).

## DNS Malware

Sophisticated malware also has emerged as a serious threat to DNS infrastructure. They are classified as Advanced Persistent Threats, and use DNS to embed themselves in the target network and stealthily communicate with external command servers to obtain malware updates, instructions, and to conduct attacks for data theft, industrial espionage and other goals.

## DNS Reflection and Amplification Attacks

As with UDP flood, DNS reflection attacks use a form of IP spoofing, changing the source address in their DNS queries to show the address of their intended target, such as a DNS root server or a top-level domain (TLD) name server operator. DNS reflection and amplification recognizes UDP as an asymmetrical protocol (small requests, large responses) and the existence of open DNS resolvers to the Internet cloud. The result is that small DNS queries reflect large UDP datagram responses to the target address in the original source datagrams. Some recent attacks have used this DDoS technique at a huge scale.

Because DNS runs over UDP and does not require a handshake, it is possible to use the protocol as a means to lock down a host or a network. Designed a specific way, sending a small query to any open DNS resolver can result in a single response containing several kilobytes or more, that are sent to the unwitting spoofed victim. (This type of response typically is sent via TCP, as UDP does not allow for more than 512 bytes in a response datagram. The resulting packet usually exceeds the MTU of the recipient's interfaces, resulting in further packet fragmentation and processing.) Open DNS resolvers may allow for launching DDoS attacks containing hundreds of gigabytes of data. Attackers may also use the EDNS0 DNS protocol extension as a means to enable larger DNS responses. Many network operators, particularly overseas, allow open DNS resolvers to run on their networks, unwittingly allowing attackers to abuse them. Many network operators do provide intelligent rate-limiting to prevent abuse, even while supporting open recursive DNS servers. Hence, issues of this type usually result from mistakes in configuration.

## DNS Cache Poisoning

With cache poisoning, attackers attempt to insert a spoofed DNS response to a DNS resolver, which then stores the response in its cache, where it lives until the TTL expires. The cache is poisoned and subsequent requests for the domain address to recursive name servers are answered with the address of a different server, presumably controlled by the attacker. So long as the fake entry resides in the DNS server cache (persistence of a cache entry is usually governed by time to live) it can result in hundreds of thousands of dangerous redirects. In such cases the URL is legitimate, but the destination servers are not. This process is often called a "pharming" attack. Web servers and mail servers are frequent targets. Other redirection attacks include DNS Changer and DNS Replay. Man-in-the-middle is another descriptive term for many redirection attacks.

## DNS Fluxing

the use of a system called a Domain Generation Algorithm botnet to perform one of the following attacks:

- **Fast Fluxing:** Forcing rapid swapping in and out of IP addresses, with extremely high frequency through changing DNS records with brief TTLs
- **Domain Fluxing:** Forcing constant changing of and allocation of multiple fully-qualified domain names (FQDNs) to a single IP address on the recursive or authoritative DNS server.

## Inside-Out Attacks

A sophisticated form of "phishing" in which an attacker is able to inject a worm or other piece of attack software onto a host machine, which thereupon captures sensitive information such as logins, and adds that data to DNS queries that can be sent from the trusted machine to an untrusted entity for collection. DDoS Security detects data leaks of this type, logs the incident, and funnels the suspect packets to a quarantine location. In a similar vein, **DNS Tunneling** uses DNS as a covert channel to avoid firewall and IPS security mechanisms. Tunneling encapsulates Inbound and outbound packets inside DNS requests and DNS responses.

## UDP DNS Flood

UDP Flood is a denial-of-service attack that uses the connectionless UDP transport protocol and attempts to send large numbers of packets to random UDP protocol ports on a remote system, or to a specific protocol. UDP flood is a reflection attack that is often used for attacking DNS servers operating on UDP port 53. UDP flooding typically uses IP spoofing, in which the sender address is faked. The purpose is to occupy so many resources on the target that it can no longer provide its services on the network.

## SYN Flood

A host sends a long stream of TCP SYN packets, frequently using a forged sender address. Because TCP regards a SYN packet as part of a legitimate connection request, the requested server starts a half-open connection by responding with a SYN ACK packet. Since the sender address is faked, the final ACK response from the sender never comes, and the half-open TCP socket closes only after a time out interval. A massive wave of SYN requests with fake senders can wipe out the connection resources of a network device, effectively locking it away from legitimate users.

## Internet Control Message Protocol (ICMP) Flood

An ICMP flood attack is also known as a ping attack in which attackers send a large number of ICMP ping packets to a DNS server repeatedly in order to hinder the server's ability to respond to other requests. It can also be an attempt to send a large number of ping packets to the broadcast IP of a subnet, otherwise known as a Smurf attack, as a basic means of amplifying an attack across more hosts than a normal ping would typically permit. These types of attacks can be dealt with by setting a policy to disallow pings to the broadcast IP on the network.

### Note

When threat protection is enabled, ICMP ping size (for IPv4 and IPv6) is limited to 16,000 bytes.

## Infoblox DNS Firewall

This section provides information about the Infoblox DNS Firewall feature that you can configure and manage on the Infoblox appliance. It includes the following topics:

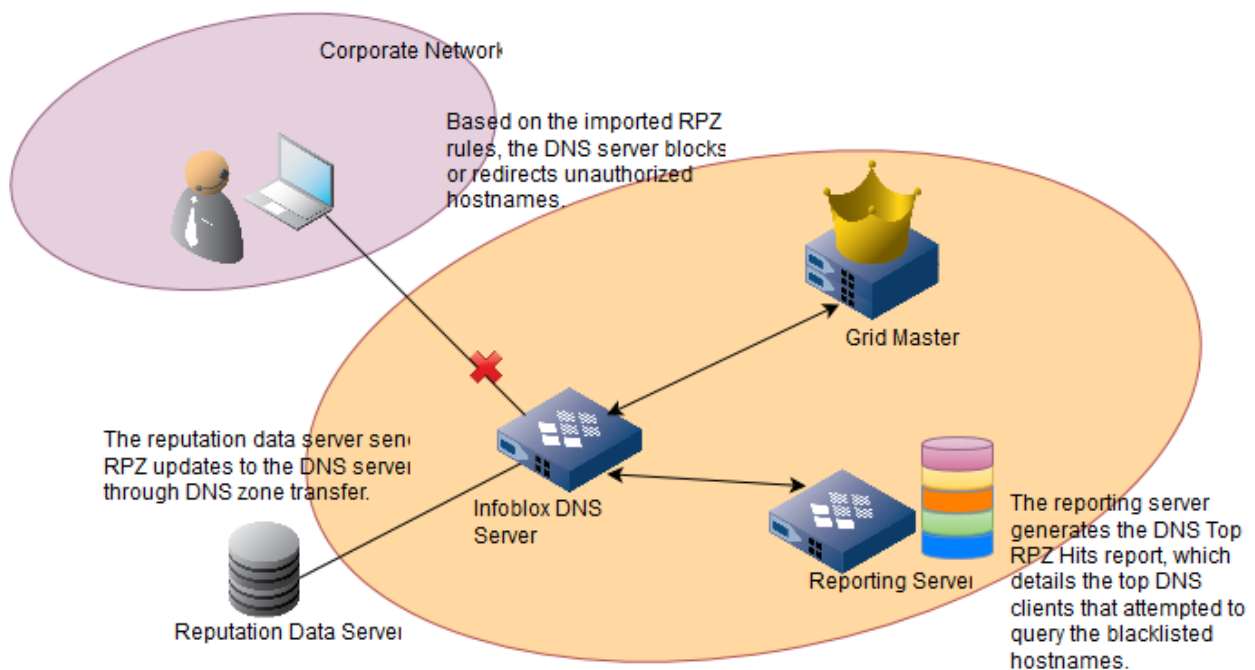
- [About Infoblox DNS Firewall](#)
- [License Requirements and Admin Permissions for RPZ](#)
- [Enabling Recursion for RPZs](#)
- [Configuring Local RPZs](#)
- [Configuring Rules for RPZs](#)
- [Configuring Infoblox Threat Intelligence Feed](#)
- [Downloading Rules for an RPZ Feed](#)
- [Testing RPZ Feed Rules](#)
- [About FireEye Integrated RPZs](#)
- [Mitigating Cyber Threats using TAXII](#)
- [Managing RPZs](#)
- [Managing RPZ Rules](#)
- [Configuring Prefix Length Limit for RPZ-IP Triggers](#)
- [Configuring Thresholds for RPZ Hit Rate](#)
- [Verifying RPZ Configuration](#)
- [Best Practices for Configuring RPZs](#)
- [Configuring Combination Threat Feeds](#)
- [Best Practices for Deploying Combination Threat Feeds](#)

## About Infoblox DNS Firewall

Infoblox DNS Firewall employs DNS RPZs (Response Policy Zones), a technology developed by ISC (Internet System Consortium) for allowing reputable sources to dynamically communicate domain name reputation so you can implement policy controls for DNS lookups.

On an Infoblox appliance, you can configure RPZs and define RPZ rules to block DNS resolution for malicious or unauthorized hostnames, or redirect clients to a walled garden by substituting responses. You can assign actions to RPZ rules. For example, abc.com can have an action of pass thru or substitute (domain) with the domain xyz.com. You can also configure a Grid member to act as a lead secondary that receives RPZ updates from external reputation sources and redistributes the updates to other Grid members. Infoblox DNS Firewall supports both IPv4 and IPv6 networks. It also facilitates the detection of malware and APTs (Advanced Persistent Threats) by integrating the NIOS appliance with a FireEye appliance. You can employ APT mitigation strategy using FireEye as an external threat detection source. An Infoblox Grid performs RPZ actions for queries that originate from external sources. The name server recursive cache on an RPZ enabled Grid member uses the address of the client from which the query originates to identify if the query is generated from an external source or an internal Grid. If the query originates from a Grid Master or a Grid member that has RPZ license installed, RPZ actions are automatically bypassed for those queries. For RPZ, Infoblox uses the ACL *infoblox-deny-rpz*, which contains a list of addresses for bypassing RPZ actions. The *infoblox-deny-rpz* list excludes Grid members that do not have an RPZ license. Note that RPZ action is performed only once for a single recursion. As illustrated in the below figure, the Infoblox DNS server receives RPZ updates, which include blacklisted hostnames and responses, from a reputation data server through a DNS zone transfer. The appliance then blocks or redirects queries and responses based on the imported RPZ rules. The reporting server can then generate the *DNS Top RPZ Hits* report that details the top DNS clients that have received redirected responses through RPZs.

*Infoblox DNS Firewall*



There are three types of RPZs:

- Local RPZ – A local RPZ is a zone that allows administrators to define multiple response policies locally. Responses sent are based on the defined rules. For information about how to configure local RPZs, see [Configuring Local RPZs](#).

- RPZ Feed – An RPZ feed receives response policies from external sources. DNS clients receive responses based on the imported rules from a reputable source, such as a commercial RPZ provider. For information about RPZ feed, see [Configuring Infoblox Threat Intelligence Feed](#).
- FireEye integrated RPZ – By integrating the NIOS appliance with the FireEye appliance, you can detect malware and APTs and take necessary actions to mitigate those threats. For information about FireEye integrated RPZ, see [About FireEye Integrated RPZs](#).



#### Note

You can configure up to a total of 32 RPZs, including local and FireEye integrated RPZs.

For more information on configuring RPZ feeds using On-Prem Firewall Service, see [On-Prem DNS Firewall Service](#).

## Setting Up Infoblox DNS Firewall

For a successful Infoblox DNS Firewall deployment to protect your endpoint devices and servers from stealthy malware and malicious hostnames, consider the guidelines described in [Best Practices for Configuring RPZs](#). To configure Infoblox DNS Firewall, complete the following tasks:

1. Install a valid RPZ license on the appliance, as described in [License Requirements and Admin Permissions for RPZ](#).  
Note to ensure that you have installed a valid DNS license on the same appliance.
2. Enable recursive queries for a DNS view, member, or Grid, as described in [Enabling Recursion for RPZs](#).  
Note to ensure that you enable recursive queries for RPZ rules to take effect.
3. Configure RPZ logging to ensure that all matching and disabled rules for all queries are logged in the syslog. You can view the syslog to ensure that the rules are set up correctly before they take effect. Ensure that you enable **rpz** in the **Logging Category** of *Grid DNS Properties* editor to log these events. For information about how to set logging categories, see [Using a Syslog Server](#).
4. You can configure a local RPZ, an RPZ feed, or a FireEye RPZ on the NIOS appliance. Complete one of the following depending on your selection:
  - On a DNS member, complete the following to create local RPZ rules:
    - i. Create an RPZ, as described in [Configuring Local RPZs](#).
    - ii. Configure rules for the local RPZ you have created, as described in [Configuring Rules for RPZs](#).
  - Optionally, complete the following to receive RPZ updates from an RPZ feed:
    - i. Configure an RPZ feed, as described in [Configuring Rules for RPZs](#). You can also configure the Infoblox DNS feed, as described in [Configuring Infoblox Threat Intelligence Feed](#). The Infoblox DNS feed is a reputable data server validated by Infoblox to provide reputation RPZ updates.
    - ii. Download rules from the RPZ feed, as described in [Downloading Rules for an RPZ Feed](#).
  - Optionally, complete the following to receive alerts from a FireEye appliance:
    - i. Create a FireEye integrated RPZ, as described in configuring fireeye RPZs, see [About FireEye Integrated RPZs](#).
    - ii. Define rules for FireEye RPZs, as described in configuring fireeye RPZs, see [About FireEye Integrated RPZs](#).
    - iii. Create FireEye admin users, as described in for fireeye integrated RPZs, see [About FireEye Integrated RPZs](#).
    - iv. Add URLs and user credentials on the FireEye appliance, as described in configuring fireeye RPZs, see [About FireEye Integrated RPZs](#).



#### Note

To apply the configured RPZ policies regardless of whether a DNS query requests DNSSEC data, configure the appliance accordingly. For more information about how to configure this, see [Applying Policies and Rules to DNS Queries that Request DNSSEC Data](#).

5. Test your RPZ configuration and verify that RPZ is functioning properly by viewing the syslog and the **Last Updated** column in the **Response Policy Zones** tab. For more information, see [Testing RPZ Feed Rules](#).

After you have set up your RPZs, RPZ feeds, and RPZ rules. You can do the following:

- Manage local RPZs such as viewing a list of RPZs, modifying, reordering, and deleting RPZs. You can also lock or unlock RPZs. For more information, see [Managing RPZs](#).
- Verify RPZs are functioning properly by viewing the syslog and the last updated RPZ. For more information, see [Managing RPZs](#).
- Manage Local RPZ rules such as viewing, modifying, and deleting RPZ rules. You can also copy and import RPZ rules. For more information, see [Managing RPZs](#).
- Generate the *DNS Top RPZ Hits* report, if you have a reporting server set up in the Grid. For more information DNS Top RPZ Hits, see [About Dashboards](#).
- Define thresholds for RPZ hit rate and configure the appliance send alerts when the RPZ hit rate exceeds the thresholds. For information, see [Configuring Thresholds for RPZ Hit Rate](#).

## License Requirements and Admin Permissions for RPZ

You must install required licenses before you can use the RPZ feature. An RPZ license is required to configure local RPZs and RPZ feeds.

This section includes the following topics:

- [For Local RPZs and RPZ Feeds](#)
- [Grid-wide licenses for RPZ](#)
- [For FireEye Integrated RPZs](#)

For more information, see [For Local RPZs and RPZ Feeds](#) below. For FireEye integrated RPZs, you must first install an RPZ license, and then a Security Ecosystem license. For more information, see [For FireEye Integrated RPZs](#) below. For all RPZ related licenses, you can install either a temporary or a permanent license on the NIOS appliance. The temporary license provides a 60-day free trial, which can be upgraded to a permanent license. After the license expires, the RPZs will remain intact, but you cannot delete existing or add new entries to it. Infoblox provides RPZ licenses that are compatible with each product model.

### For Local RPZs and RPZ Feeds

Before you install an RPZ license, ensure that the following are completed:

- The entire Grid is running NIOS 6.6 or later.
- Grid members are properly configured and DNS is enabled on the members.



#### Note

Install RPZ licenses only on Infoblox members that have DNS recursion enabled.

Superusers can configure RPZs and RPZ rules by default. You can also assign global permissions for all RPZs and RPZ rules to specific admin groups and roles. For more information about administrative permissions for zones, see [Administrative Permissions for DNS Resources](#).

## Grid-wide licenses for RPZ

Infoblox offers a Grid-wide **BloxOneThreat Defense Business On-Premises, BloxOneThreat Defense Business Cloud, and BloxOneThreat Defense Advanced** additional services for DNS Firewall feature. These are subscription licenses that you may have to purchase from Infoblox. When you purchase these services, you can use the BloxOne threat subscription data on any appliance that can run DNS Firewall across the Infoblox Grid which you have configured. You can buy these additional services if you do want to use Infoblox threat data and instead use your own or an external threat feed. With these additional services, you get to use a separate license key for each Grid that you use.

When you configure Grid-wide RPZ license for each Grid, RPZ rules are not applied for queries that originate from the other RPZ member(s) of the Grid. When you use multiple Infoblox Grids, you can use the same TSIG key to configure feed zones across all Grids to synchronize feed subscriptions. Note that these services are valid for NIOS versions 8.0 and above. For more information about managing Grid-wide licenses, see [Managing Licenses](#).

For a thorough understanding of the RPZ process using various NIOS versions, consider a scenario where you have configured four members on the Grid: GM(RPZ license), M1(RPZ license), M2(DNS license), M3(RPZ license), and RM1 (Reporting Member).

- In NIOS version 7.3, all members with RPZ licenses are added to the ACL list:

```
infoblox-deny-rpz { localhost; GM_IP, Member1_IP; Member3_IP; }
```

- In NIOS version 8.x:
  - When you install a Grid Wide RPZ license, all members with DNS licenses are added to the ACL list:

```
infoblox-deny-rpz { localhost; GM_IP, Member1_IP; Member2_IP; Member3_IP; }
```

- When you install a Grid Wide RPZ license, but not a member level RPZ license, NIOS adds all the members with DNS licenses to the ACL list:

```
infoblox-deny-rpz { localhost; GM_IP, Member1_IP; Member2_IP; Member3_IP; }
```

- When you do not install a Grid-wide RPZ license but install member-level RPZ licenses, NIOS adds all members that have RPZ licenses to the ACL list and enables the **Apply RPZ rules only on this member if possible** check box for that member. You can clear the check box.

```
infoblox-deny-rpz { localhost; GM_IP, Member1_IP; Member2_IP; Member3_IP; }
```

### Note

Infoblox suggests that when you upgrade to NIOS 8.0 and later versions, select the check box **ApplyRPZrulesonlyonthismemberifpossible** only if there is an existing RPZ license for the Grid member to ensure that there is no change in behavior due to an upgrade.

Note the following about Grid-wide licenses for RPZ:

- NIOS displays **rpz** in the Grid-wide and Member tabs when RPZ license is installed.
- When you enable a temporary license for RPZ, it is enabled at the Grid level and listed under the **Grid Wide** tab. You do not have to set the license at the member level.
- NIOS performs a pre-provision check for an RPZ member who joins the Grid to verify if a Grid-wide license is installed for the RPZ member. If the Grid-wide license is found, NIOS allows the member to join the Grid.
- If an RPZ license expires, Feed Zone stops receiving feed updates after the grace period. However, the RPZ feature remains until it expires. This grace period TTL is configured in BloxOne Plus/Advanced services.

- After you remove the Grid-wide license for RPZ, this feature still keeps functioning for members that do not have a Grid-wide license until you restart the service. When you remove the RPZ license that is associated with a specific member, NIOS restarts the service automatically and disables RPZ for the respective member.
- NIOS checks for a Grid-wide license of a pre-provisioned type when a member first joins the Grid. For example, if a member is pre-provisioned for RPZ and a Grid-wide RPZ license is installed, then NIOS allows the member to join the Grid, even though the member does not have a valid RPZ license.

### Customizing Threat Feeds on an RPZ Member

To configure an RPZ member to use your own or external threat feeds instead of RPZ feeds, complete the following:

1. From the **Data Management** tab, select **DNS**, and then click **Members**.
2. Select a member and click **Edit** from the Toolbar.
3. In the *Member DNS Properties* editor, click **General** tab -> **Advanced** tab.
  - **Apply RPZ rules only on this member if possible:** Select this check box if the forwarders must not apply RPZ rules to the responses that is returned to the other member, when this RPZ member queries other Grid member details.
4. Save the configuration.

### For FireEye Integrated RPZs

You can enable FireEye integrated RPZs on the appliances that have both the RPZ and Security Ecosystem licenses installed. Note that you must install an RPZ license prior to installing the Security Ecosystem license.

NIOS appliance creates a new group, **fireeye-group**, when you add the first FireEye zone. The FireEye admin group is read-only and you cannot assign permissions to it. It will not have any superuser privileges and you cannot modify or delete this group. You can add users to the **fireeye-group** admin group, and FireEye users can only send alerts to the NIOS appliance. They cannot access the Infoblox GUI, CLI, API, or RESTful API. These users are authenticated based on the usernames and passwords you configure in the FireEye admin group. Only admin users who belong to the FireEye admin group can publish FireEye alerts. Other admin users cannot do so. For information about how to configure the FireEye appliance, see [Configuring the FireEye Appliance](#).

#### Note

The **fireeye-group** is created automatically. Infoblox recommends that you do not add a group with the same name. In addition, The "force password change at next login" feature does not apply to admin users in the **fireeye-group**. These users will not be prompted to change their passwords at the next login. Their original passwords continue to work. For more information managing passwords, see [Creating Local Admins](#).

To add users to the **fireeye-group**, complete the following:

1. From the **Administration** tab, select **Administrators**, and then click **Admins**.
2. Click **Add** and enter the usernames and passwords. For more information on how to add users to an admin group, see [Creating Local Admins](#). Select **fireeye-group** for the admin group and add users to this group.

#### Note

Ensure that you save the usernames and passwords. You must use these credentials when configuring FireEye alerts to enable the alerts to be received by NIOS.

### Uninstalling the Security Ecosystem License

When you uninstall the Security Ecosystem license, new FireEye alerts will not be processed. However, the FireEye integrated RPZs and the rules in those zones will not be deleted. Note the following when you uninstall the Security Ecosystem license:

- New FireEye alerts will not be processed



- FireEye RPZ zones that were created before uninstalling the license will remain
- You cannot create new FireEye RPZ zones
- RPZ rules created from the alert will remain
- Note that if the RPZ and Security Ecosystem licenses are installed, then you must first remove the Security Ecosystem license to remove the RPZ license.
- The fireeye-group and the FireEye zones will remain even after you delete the Security Ecosystem license.

## Enabling Recursion for RPZs

For RPZ rules to function properly, you must enable DNS recursion. You can enable DNS recursion at the Grid, member, or DNS view level. To enable recursion:

- For the Grid or member, see [Enabling Recursive Queries](#).
- For a DNS view, see [Managing Recursive DNS Views](#).

## Configuring RPZs for All Recursive Servers

When you configure a local or FireEye integrated RPZ, you must define an internal primary name server. The primary name server can be either recursive or non-recursive, depending on its usage. When you configure an RPZ feed, you must define an external primary name server. You can associate a name server or a name server group with the local RPZ, RPZ feed, or FireEye RPZ. You can also configure RPZs and RPZ feeds for all recursive servers in the Grid. A local RPZ can have one or more secondary name servers associated with it. For an RPZ feed, you must create an external primary name server.

To configure a local RPZ, or RPZ feed, or FireEye RPZ for all recursive servers, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then click the **Add** icon.
2. Enter the *Response Policy Zone* details and click **Next** to associate an RPZ with at least one name server. For information about creating a local RPZ, and creating an RPZ feed, see [Configuring Local RPZs](#). For information about creating a FireEye integrated RPZ, see [Configuring FireEye RPZs](#).
3. Select **All Recursive Name Servers** from the list to add all the recursive name servers in the Grid as the secondary name servers for the corresponding zone.
  - Save the configuration and click **Next** to define extensible attributes. Click **Restart** if it appears at the top of the screen. For information about extensible attributes, see [Managing Extensible Attributes](#).

## Enabling and Disabling RPZ Query Name Recursion

In previous NIOS releases, RPZ query name recursion was enabled by default. The DNS recursive name server performed RPZ recursive lookups for the fully qualified domain name that was part of an RPZ. Starting with NIOS 7.1.0, RPZ query name recursion is disabled by default. When RPZ query name recursion is disabled, the DNS recursive name server sends responses for the domains being queried, without forwarding queries to the authoritative name servers. This can speed up recursive RPZ lookups by eliminating unnecessary recursions for domains that are known to be malicious, possibly caused by internal DDoS attacks on the recursive server.

You can enable RPZ query name recursion by selecting the **Enable RPZ query name recursion (qname-wait-recurse)** checkbox. When you select this checkbox, the appliance performs RPZ query name recursions. You can configure this at the Grid, member, and DNS view levels.

### Note

RPZ query name recursion is disabled by default. The **EnableRPZquerynamerecursion(qname-wait-recurse)** checkbox is deselected for all new installations and upgrades. You can select this checkbox to enable RPZ query name recursion.

To enable or disable RPZ query name recursion:

1. From the **Data Management** tab, select the **DNS** tab, and then click **Grid DNS Properties** from the Toolbar.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *Grid\_member* checkbox, and then click the Edit icon.



or

From the **Data Management** tab, select the **DNS** tab -> **Zones** tab -> *dns\_view* checkbox, and then click the Edit icon.

2. In the *Grid DNS Properties*, *Member DNS Properties*, or *DNS View* editor, click the **General** tab -> **Advanced** tab and complete the following.
  - **Enable RPZ query name recursion (qname-wait-recurse)**: This checkbox is deselected by default, meaning RPZ query name recursion is disabled. When RPZ query name recursion is disabled, the DNS recursive name server sends responses for the domains being queried, without forwarding queries to the authoritative name servers. When you select this checkbox, the DNS recursive name server performs recursive lookups for the fully qualified domain names that are part of an RPZ. To override the value inherited from the Grid, click **Override**. To retain the same value as the Grid, click **Inherit**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

Following are sample configuration details in the *named.conf* file when the **Enable RPZ query name recursion (qname-wait-recurse)** checkbox is deselected:

```
response-policy {  
    zone "local.com" policy Given;# priority 0  
    zone "rpz.net" policy Given;# priority 1  
    zone "example.com" policy Given;# priority 2  
    } qname-wait-recurse no ;  
    dnssec-enable yes;  
    dnssec-validation yes;  
    dnssec-accept-expired no;  
    filter-aaaa-on-v4 no;  
    zone "." in {  
        type hint;  
        file "named.cache._default";  
    };
```

## Configuring Local RPZs

You can define local RPZs to match responses for recursive queries. Each RPZ can have various rules associated with it. The response of a recursive query is modified if it matches any of the RPZ rules. The responses are first matched with the RPZ rules, and if there is a match, the rule defined at the RPZ level override is used. When creating a new RPZ zone, you can associate this zone with a threat severity level. The RPZ syslog messages provide information about threat severity level of an RPZ zone associated with the matched RPZ rule. To view threat details, you can drill down to the syslog messages. For more information about viewing RPZ in the syslog, [Verifying RPZ Configuration](#). You can create multiple local RPZs and define multiple rules for a local RPZ. Note that override depends on the order of the zones. The zones on top will override the zones below. You can change the order of the RPZs. For more information reordering RPZs, see [Managing RPZs](#). You can also configure FireEye integrated RPZs on the NIOS appliance to detect persistent threats and malwares. The NIOS appliance considers the FireEye integrated RPZ as a local RPZ. For more information, see [About FireEye Integrated RPZs](#).



#### Note

When using IDN (Internationalized Domain Name) in a local RPZ or RPZ feed, you must manually convert the IDN to punycode. For information about IDN, see [Support for Internationalized Domain Names](#).

To configure local RPZs:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then click the *Add* icon.
2. When you click the *Add* icon, either the *Add Response Policy Zone Wizard* or the *Add DNS View* wizard is displayed based on the following:
  - When you click the *Add* icon, the *Add Response Policy Zone Wizard* is displayed if you have not created additional DNS views and only have the default view.
3. If you have configured multiple DNS views, you must drill-down to the corresponding view to assign a local RPZ. Click the *Add* icon and the *Add Response Policy Zone Wizard* is displayed. To create a new DNS view for your local RPZ, click the *Add* icon and complete the details in the *Add DNS View* wizard. For information about adding a DNS view, see [Adding a DNS View](#).
  - In the *Add Response Policy Zone Wizard*, select **Add Local Response Policy Zone**, click **Next** and specify the following:
    - **Name:** Enter the name of the local RPZ. It can be a combination of alphanumeric characters. You can enter up to 256 characters.
    - **DNS View:** The name of the view that you have selected is displayed by default. You can select a view from the drop-down list to associate it with the local RPZ.
  - Note that the local RPZ must have a primary Grid name server before you can configure it.
  - **Policy Override:** Select a value from the drop-down list. You can override the policy actions that are specified in the rule level.
    - **Log Only (Disabled)** – Select this if you want to disable an RPZ rewrite using rules in the RPZ. If the response to the recursive query matches any RPZ rule, then the rule is logged, but the response will not be altered. Note that this option will not override RPZ rules in other RPZ zones, if they take precedence. Select this option to preview the rules in the syslog before they take effect.
    - **None (Given)** – Select this if you want to use the policy from the rule level.
    - **Block (No Data)** – Select this to send a response that contains no data in it.
    - **Block (No Such Domain)** – Select this if you want the user to receive a DNS response that indicates there is no domain. All the policy actions in an RPZ are replaced with a NXDOMAIN block.
    - **Pass thru** – Select this if you want to send an actual response without modification. All the policy actions in an RPZ are replaced with the passthru action.
    - **Substitute (Domain Name)** – Select this if you want to replace all the policy actions in an RPZ with the specified substitution action.
    - **Domain Name:** This appears only when you select **Substitute (Domain Name)** from the **Policy Override** list. Enter the domain name that you want the client to receive instead of the actual domain name, which is malicious or unauthorized.
  - **Severity:** Select the threat severity level for the RPZ zone. The threat severity you select here determines the severity for the RPZ zone. Select Critical, Major, Warning, or Informational. The default threat severity level is Major. Note that each of these levels is represented by a number in the syslog (8 being Critical and 4 being Informational). When you upgrade to NIOS 7.0.0, the appliance automatically updates the threat severity level to Informational (displayed as 4 in the syslog) for existing RPZ zones. For information about viewing RPZ in the syslog RPZ and severity levels, see [Verifying RPZ Configuration](#).
  - **Comment:** Optionally, enter additional information about the local RPZ.
  - **Disable:** Select the checkbox to disable a local RPZ without deleting its configuration. Clear the checkbox to enable the local RPZ. For information, see [Enabling and Disabling Zones](#). Note that disabling a local RPZ may take a longer time to complete depending on the size of the data.
  - **Lock:** Select the checkbox to lock the zone so that you can make changes to it and prevent others from making conflicting changes. As described in locking and unlocking RPZs, see [Managing RPZs](#).
4. Click **Next** to associate the local RPZ with at least one primary name server:
  - Define the name servers for the local RPZ. A Grid name server must be recursive when primary Grid name server is used as an RPZ source. A local RPZ may or may not have a recursive server. For

example, there could be a Grid that has only primary Grid name server for a local RPZ to act as an RPZ source for an external set of name servers. A local RPZ must have only one primary Grid name server and it can have one or more secondary Grid name servers. When you select **All Recursive Name Servers** from the list, all the recursive name servers in the Grid are added as secondary servers for the zone. For information on specifying primary or secondary name servers, see [Assigning Zone Authority to Name Servers](#). For information on specifying name server groups, see [Using Name Server Groups](#). For information about configuring RPZs for all recursive servers, see [Enabling Recursion for RPZs](#).

5. Save the configuration and click **Next** to define extensible attributes. Click **Restart** if it appears at the top of the screen. For information, see [Managing Extensible Attributes](#).



#### Note

You cannot convert a local RPZ to an RPZ feed or vice versa.

## Configuring Rules for RPZs

You can define different RPZ rules to block DNS resolution for malicious or unauthorized hostnames or redirect clients to a walled garden by substituting responses. Depending on the nature of the rule and its usage, each rule is designed to match a hostname, domain name, or IP address, specification or pattern, and an associated action.

These rules are applicable to local RPZs, including FireEye integrated RPZs, except for the RPZ client IP address or network rules which are not applicable for FireEye integrated RPZs. For RPZ feeds, rules are imported from external servers. You cannot change the content of an RPZ feed, but you can override the actions in an RPZ feed.

The RPZ rules are triggered based on the order of the RPZ zones that you have configured. When you configure one or more RPZ rules with the same FQDNs or IP addresses in different RPZ zones, then the RPZ rules in the top-level RPZ zone are triggered first.

This section includes the following topics:

- [Managing Passthru Rules](#)
- [Managing Block \(No Such Domain\) Rules](#)
- [Managing Block \(No Data\) Rules](#)
- [Managing Substitute \(Domain Name\) Rules](#)
- [Managing Substitute \(Record\) Rules](#)

To configure RPZ rules:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, click *DNS\_View* -> *Zone* and then click **Add** -> select a **Rule**.
2. The rules are classified as follows:
  - **Passthru Rule**
  - **Block (No Such Domain) Rule**
  - **Block (No Data) Rule**
  - **Substitute (Domain Name) Rule**
  - **Substitute (Record) Rule**
3. Complete the details in the corresponding editor.
4. Save the configuration and click **Next** to define extensible attributes.

You cannot define the above rules for an RPZ feed. An RPZ feed uses rules defined by external servers. When you click on an RPZ feed, the appliance displays a dialog box that provides various options to export the rules of the configured external servers in .CSV format.

## Managing Passthru Rules

You can define passthru rules if you do not want to modify the actual responses of the recursive queries. The response received for a query is not modified, if there is a matching passthru rule and the actual response is forwarded to the user.

## Adding Passthru Rules for Domain Names

To define passthru rules for domains:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Rule** -> select **Passthru Rule** -> **Passthru Domain Name Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Rule** -> select **Passthru Rule** -> **Passthru Domain Name Rule**.
2. The following fields are displayed in the *Add a Passthru Domain Name Rule* wizard:
  - **Name:** Enter the domain name for which you want to define the passthru rule. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Policy:** Displays the selected policy.
  - **Comment:** Optionally, enter additional information.
  - **Disable:** Clear the checkbox to enable the passthru rule. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
4. Save the configuration.

## Adding Passthru Rules for IP Addresses or Networks

To define passthru rules for IP addresses or networks:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Rule** -> select **Passthru Rule** -> **Passthru IP Address Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Rule** -> select **Passthru Rule** -> **Passthru IP Address Rule**.
2. The following fields are displayed in the *Add a Passthru IP Address Rule* wizard:
  - **IP Address or Network:** Enter the IP address or specify the address in CIDR format for which you want to define the passthru rule. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Policy:** Displays the selected policy.
  - **Comment:** Optionally, enter additional information.
  - **Disable:** Clear the checkbox to enable the passthru rule. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
4. Save the configuration.

## Adding Passthru Rules for Client IP Addresses or Networks

You can define a passthru rule for a client IP address or network, if you do not want to modify the response to a query from a specific client IP address or network and forward the actual response to the client.

To define passthru rules for the client IP addresses or networks:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Rule** -> Select **Passthru Rule** -> **Passthru Client IP Address Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Rule** -> select **Passthru Rule** -> **Passthru Client IP Address Rule**.
2. The following fields are displayed in the *Add a Passthru Client IP Address Rule* wizard:
  - **Client IP Address or Network:** Enter the client IP address or specify the client address in CIDR format for which you want to define the passthru rule. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Policy:** Displays the selected policy.
  - **Comment:** Optionally, enter additional information.
  - **Disable:** Clear the checkbox to enable the passthru rule. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
4. Save the configuration.

## Managing Block (No Such Domain) Rules

You can define rules to block certain domain names, IP addresses or networks, or client IP addresses or networks. When you choose this option to block a domain name, the query name is matched with the RPZ rule. If the query name matches the RPZ rule, the DNS client receives a DNS response that indicates the domain does not exist. When you block an IP address or network using this option, the A and AAAA records are matched with the RPZ rule. If the records match an RPZ rule, the DNS client receives a DNS response that indicates the domain does not exist. When you choose this option to block a specific client IP address or network, the IP address or network of a client querying the DNS server is matched with the RPZ rule. If the IP address or the network of the client matches the RPZ rule, the DNS client receives a DNS response that indicates the domain does not exist.

### Defining Block (No Such Domain) Rules for Domain Names

To define block rules for domains:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Block (No Such Domain) Rule** -> **Block Domain Name (No Such Domain) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Block (No Such Domain) Rule** -> **Block Domain Name (No Such Domain) Rule**.
2. The following fields are displayed in the *Add a Block Domain Name (No Such Domain) Rule* wizard:
  - **Name:** Enter the domain name which you want to be blocked from being resolved by the DNS. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Policy:** Displays the selected policy.
  - **Comment:** Optionally, enter additional information.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
4. Save the configuration.

### Defining Block (No Such Domain) Rules for IP Addresses or Networks

To define block rules for IP addresses or networks:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Block (No Such Domain) Rule** -> **Block IP Address (No Such Domain) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Block (No Such Domain) Rule** -> **Block IP Address (No Such Domain) Rule**.
2. The following fields are displayed in the *Add a Block IP Address (No Such Domain) Rule* wizard:
  - **IP Address or Network:** Enter the IP address or specify the address in CIDR format which you want to block. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Policy:** Displays the selected policy.
  - **Comment:** Optionally, enter additional information.
  - **Disable:** Clear the checkbox to enable the block rule. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
4. Save the configuration.

### Defining Block (No Such Domain) Rules for Client IP Addresses or Networks

To define block rules for client IP addresses or networks:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Block (No Such Domain) Rule** -> **Block Client IP Address (No Such Domain) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone.

Expand the Toolbar, click **Add** -> select **Block (No Such Domain) Rule** -> **Block Client IP Address (No Such Domain) Rule**.

- The following fields are displayed in the *Add a Block Client IP Address (No Such Domain) Rule* wizard:
  - Client IP Address or Network:** Enter the client IP address or specify the client address in CIDR format which you want to block. Click **Select Zone** to select a different zone.
  - DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - Policy:** Displays the selected policy.
  - Comment:** Optionally, enter additional information.
  - Disable:** Clear the checkbox to enable the block rule. Select the checkbox to disable it.
- Click **Next** to define extensible attributes.
- Save the configuration.

## Managing Block (No Data) Rules

You can define rules to block certain domain names, IP addresses or networks, or client IP addresses or networks. When you choose this option to block a domain name, the query name is matched with the RPZ rule. If the query name matches the RPZ rule, the DNS client receives a DNS response that indicates there is no data for the requested record type.

When you block an IP address or network using this option, the A and AAAA records are matched with the RPZ rules. If the records match an RPZ rule, the DNS client receives a DNS response that indicates there is no data for the requested record type.

When you choose this option to block a specific client IP address or network, the IP address or network of a client querying the DNS server is matched with the RPZ rule. If the IP address or the network of the client matches the RPZ rule, the DNS client receives a DNS response that indicates there is no data for the requested record type.

## Defining Block (No Data) Rules for Domain Names

To define block rules for domains:

- From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Block (No Data) Rule** -> **Block Domain Name (No Data) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Block (No Data) Rule** -> **Block Domain Name (No Data) Rule**.
- The following fields are displayed in the *Add a Block Domain Name (No Data) Rule* wizard:
  - Name:** Enter the domain name which you want to block. Click **Select Zone** to select a different zone.
  - DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - Policy:** Displays the selected policy.
  - Comment:** Optionally, enter additional information.
  - Disable:** Clear the checkbox to enable the block rule. Select the checkbox to disable it.
- Click **Next** to define extensible attributes.
- Save the configuration.

## Defining Block (No Data) Rules for IP Addresses or Networks

To define block rules for IP addresses or networks:

- From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Block (No Data) Rule** -> **Block IP address (No Data) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Block (No Data) Rule** -> **Block IP address (No Data) Rule**.
- The following fields are displayed in the *Add a Block IP Address (No Data) Rule* wizard:
  - IP Address or Network:** Enter the IP address or specify the address in CIDR format which you want to block. Click **Select Zone** to select a different zone.
  - DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - Policy:** Displays the selected policy.
  - Comment:** Optionally, enter additional information.

- **Disable:** Clear the checkbox to enable the block rule. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
  4. Save the configuration.

## Defining Block (No Data) Rules for Client IP Addresses or Networks

To define block rules for client IP addresses or networks:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Block (No Data) Rule** -> **Block Client IP address (No Data) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Block (No Data) Rule** -> **Block Client IP address (No Data) Rule**.
2. The following fields are displayed in the *Add a Block Client IP Address (No Data) Rule* wizard:
  - **Client IP Address or Network:** Enter the client IP address or specify the client address in CIDR format which you want to block. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Policy:** Displays the selected policy.
  - **Comment:** Optionally, enter additional information.
  - **Disable:** Clear the checkbox to enable the block rule. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
4. Save the configuration.

## Managing Substitute (Domain Name) Rules

You can define an alternate IP address or a domain name to redirect a domain name or an IP address, which is malicious or unauthorized. When the response to the client query matches an RPZ rule, the actual domain name or IP address is substituted with the alternative domain name or IP address. The client will receive the substituted value instead of the actual response.

### Note

The domain name and substitute name for which you want to define a substitute rule are not case-sensitive. For example, if a domain name is "*corpxyz.com*" and you want to substitute it with "*corpxyz.com*" or "*corpxyz.com*," the substitute rule you define becomes a passthru rule because no substitution will occur since "*corpxyz.com*" is the same as "*corpxyz.com*" and "*corpxyz.com*." Grid Manager displays such substitute rule as a passthru rule.

## Defining Substitute Domain Name (Based on Domain Name) Rules

To define substitutes for domain names:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Substitute (Domain Name) Rule** -> **Substitute Domain Name (Domain Name) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Substitute (Domain Name) Rule** -> **Substitute Domain Name (Domain Name) Rule**.
2. The following fields are displayed in the *Add a Substitute (Domain Name) Rule* wizard:
  - **Name:** Enter the domain name for which you want to define a substitute. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Policy:** Displays the selected policy.
  - **Substituted Name:** Enter an alternative domain name that has to be substituted with the actual domain name. Click **Select Zone** to select a different zone.
  - **Comment:** Optionally, enter additional information.
  - **Disable:** Clear the checkbox to enable the substitute rule. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.

4. Save the configuration.

## Defining Substitute Domain Name (Based on IP address) Rules

To define substitutes for IP addresses:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Substitute (Domain Name) Rule** -> **Substitute Domain Name (IP Address) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Substitute (Domain Name) Rule** -> **Substitute Domain Name (IP Address) Rule**.
2. The following fields are displayed in the *Add a Substitute Domain Name (IP Address) Rule* wizard:
  - **IP address or Network:** Enter the IP address or network for which you want to define a substitute. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Policy:** Displays the selected policy.
  - **Substituted Name:** Enter an alternative domain name or IP address that has to be substituted with the actual IP address. Click **Select Zone** to select a different zone.
  - **Comment:** Optionally, enter additional information.
  - **Disable:** Clear the checkbox to enable the substitute rule. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
4. Save the configuration.

## Defining Substitute Domain Name (Based on Client IP address) Rules

You can define a substitute domain name rule for a client IP address if you want to substitute the actual response to a query from the DNS client with an alternate domain name or IP address. When the IP address of the client querying a DNS server matches the RPZ rule, the actual response is substituted with the alternative domain name or IP address specified in the RPZ rule.

To define substitute domain name rule for client IP addresses:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Substitute (Domain Name) Rule** -> **Substitute Domain Name (Client IP Address) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Substitute (Domain Name) Rule** -> **Substitute Domain Name (Client IP Address) Rule**.
2. The following fields are displayed in the *Add a Substitute Domain Name (Client IP Address) Rule* wizard:
  - **Client IP address or Network:** Enter the client IP address or client network for which you want to define a substitute domain name (client IP address) rule. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Policy:** Displays the selected policy.
  - **Substituted Name:** Enter an alternative domain name or IP address that replaces the actual DNS response. Click **Select Zone** to select a different zone.
  - **Comment:** Optionally, enter additional information.
  - **Disable:** Clear the checkbox to enable the substitute rule. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
4. Save the configuration.

## Managing Substitute (Record) Rules

You can define a substitute record for a domain name, which is considered malicious. You can define substitutes for the following in a zone:

- **A records:** For information about defining substitutes for A records, see the following section [Defining Substitutes Rules for A Records](#).



- **AAAA records:** For information about defining substitutes for AAAA records, see Defining Substitute Rules for AAAA Records below.
- **MX records:** For information about defining substitutes for MX records, see Defining Substitute Rules for MX Records below.
- **NAPTR records:** For information about defining substitutes for NAPTR records, see Defining Substitute Rules for NAPTR Records below.
- **PTR records:** For information about defining substitutes for PTR records, see Defining Substitute Rules for PTR Records below.
- **SRV records:** For information about defining substitutes for SRV records, see Defining Substitute Rules for SRV Records below.
- **TXT records:** For information about defining substitutes for TXT records, see Defining Substitute Rules for TXT Records below.
- **IPv4 address:** For information about defining substitutes for IPv4 addresses, see Defining Substitute Rules for IPv4 Addresses or Networks below.
- **IPv6 address:** For information about defining substitutes for IPv6 addresses, see Defining Substitute Rules for IPv6 Addresses or Networks below.

You can define a substitute for a certain owner name and record type. When you substitute a record for a certain owner name and record type, then responses to queries for that owner name and type are modified to contain the substituted value(s).

#### Defining Substitutes Rules for A Records

An RPZ A (address) record maps a domain name to a substitute IPv4 address. To define a specific name-to-address mapping, add an A record to a previously defined RPZ.

To define substitute rules for A records:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (A Record) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Substitute (Record) Rule** -> **Substitute (A Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (A Record) Rule* wizard:
  - **Name:** Enter the domain name that you want to map to an IP address. The name that you specify, irrespective of the RPZ name, is used to determine a match for the RPZ rule. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **IP Address:** Enter the IPv4 address to which you want the domain name to map.
  - **Comment:** Optionally, enter additional information about the A record.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
4. Save the configuration.

#### Defining Substitute Rules for AAAA Records

An RPZ AAAA (address) record maps a domain name to a substitute IPv6 address. To define a specific name-to-address mapping, add an RPZ AAAA record to a previously defined RPZ.

To define substitute rules for AAAA records:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (AAAA Record) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Substitute (Record) Rule** -> **Substitute (AAAA Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (AAAA Record) Rule* wizard:
  - **Name:** Enter the domain name that you want to map to an IP address. The name that you specify, irrespective of the RPZ name, is used to determine a match with the RPZ rule. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.

- **IP Address:** Enter the IPv6 address to which you want the domain name to map.
  - **Comment:** Optionally, enter additional information about the AAAA record.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
  4. Save the configuration.

### Defining Substitute Rules for MX Records

An RPZ MX (mail exchanger) record maps a domain name to a mail exchanger. A mail exchanger is a server that either delivers or forwards mail. A wildcard MX record applies to an RPZ and all its subdomains of the owner name.

To define substitute rules for MX records:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (MX Record) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Substitute (Record) Rule** -> **Substitute (MX Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (MX Record) Rule* wizard:
  - **Mail Destination:** Enter the owner name of the MX record you want to substitute.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Host Name Policy:** Displays the hostname policy of the selected zone. Ensure that the hostname you enter complies with the hostname restriction policy defined for the zone.
  - **Mail Exchanger:** Enter the fully qualified domain name of the mail exchanger.
  - **Preference:** Select an integer from 10 to 100. The preference determines the order in which a client attempts to contact the target mail exchanger.
  - **Comment:** Optionally, enter additional information about the MX record.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
4. Save the configuration.

### Defining Substitute Rules for NAPTR Records

A DNS NAPTR object represents a Naming Authority Pointer (NAPTR) resource record. This resource record specifies a regular expression-based rewrite rule that, when applied to an existing string, produces a new RPZ name or URI.

To define substitute rules for NAPTR records:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (NAPTR Record) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Substitute (Record) Rule** -> **Substitute (NAPTR Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (NAPTR Record) Rule* wizard:
  - **Domain:** Enter the domain name to which this resource record refers. Make sure that you enter a valid FQDN. Example: test.com, foo.com, etc. The name that you specify, irrespective of the RPZ name, is used to determine a match with the RPZ rule. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Service:** Select a service from the drop-down list. This field specifies the service and protocol that are used to communicate with the host at the domain name.
  - **Flags:** The Flag field indicates whether the current lookup is terminal; that is, the current NAPTR record is the last NAPTR record for the lookup. It also provides information about the next step in the lookup process. The flags that are currently used are:
    - **U:** Indicates that the output maps to a URI (Uniform Record Identifier).
    - **S:** Indicates that the output is a domain name that has at least one SRV record. The DNS client must then send a query for the SRV record of the resulting domain name.
    - **A:** Indicates that the output is a domain name that has at least one A or AAAA record. The DNS client must then send a query for the A or AAAA record of the resulting domain name.
    - **P:** Indicates that the protocol specified in the Service field defines the next step or phase.
  - **Order:** Select an Integer from 10 to 100, or enter a value from 0 to 65535. This value indicates the order in which the NAPTR records must be processed. It processes the record with the lowest value first.

- **Preference:** Select an Integer from 10 to 100, or enter a value from 0 to 65535. Similar to the Preference field in MX records, this value indicates which NAPTR record the DNS client should process first when the records have the same Order values. It processes the record with the lowest value first.
  - **REGEX:** The regular expression that is used to rewrite the original string from the client into a domain name. RFC 2915 specifies the syntax of the regular expression. Note that the appliance validates the regular expression syntax between the first and second delimiter against the Python re module, which is not 100% compatible with POSIX Extended Regular Expression as specified in the RFC. For information about the Python re module, refer to <http://docs.python.org/release/2.5.1/lib/module-re.html>.
  - **Replacement:** This specifies the domain name for the next lookup. The default is a dot (.), which indicates that the regular expression in the REGEX field provides the replacement value. Alternatively, you can enter the replacement value in FQDN format.
  - **Comment:** Optionally, enter additional information about the NAPTR record.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
  4. Save the configuration.

### Defining Substitute Rules for PTR Records

In a forward-mapping zone, a PTR (pointer) record maps a domain name to another domain name. In an RPZ, a PTR (pointer) record maps an address to a domain name. To define a specific address-to-name mapping, add an RPZ PTR record to a previously defined RPZ.

To define substitute rules for PTR records:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (PTR Record) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Substitute (Record) Rule** -> **Substitute (PTR Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (PTR Record) Rule* wizard: You can select either **Name** or **IP address** from the drop-down list.
  - **Name:** Enter a domain name for which you want to create a pointer to another domain. The name that you specify, irrespective of the RPZ name, is used to determine a match with the RPZ rule. Click **Select Zone** to select a different zone. The name should be in the following format for RPZ:

ipaddress.in-addr.arpa.

Note that the IP address should be in the reverse format. For example, if the IP address is 10.2.1.4, then the name format for RPZ is 4.1.2.10.in-addr.arpa. The following fields are displayed when you select **Name** from the drop-down list:

- **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Domain Name:** Enter the domain name to which you want the PTR record to point. Make sure that you enter a valid FQDN. Example: test.com, foo.com, etc.
- **IP Address:** Enter an IP address for which you want to create a pointer to a domain. The following fields are displayed when you select **IP Address** from the drop-down list:
    - **Zone:** Displays the RPZ you have selected. Click **Select Zone** to select a different zone.
    - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
    - **Domain Name:** Enter the domain name to which you want the PTR record to point. Make sure that you enter a valid FQDN. Example: test.com, foo.com, etc.
  - **Comment:** Optionally, enter additional information about the PTR record.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
  4. Save the configuration.

### Defining Substitute Rules for SRV Records

A DNS RPZ SRV object represents an SRV resource record, which is also known as a service record. You can define a substitute for an SRV record. When the response to a user's query matches with an RPZ rule, then the combination of actual service, protocol, domain name and the zone is substituted with a combination of priority, weight, port and target

details that you specify.

To define substitute rules for SRV records:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (SRV Record) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Substitute (Record) Rule** -> **Substitute (SRV Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (SRV Record) Rule* wizard:
  - **Display input as:** Select the format in which you want the SRV record to be displayed. When you select RFC 2782 format, the appliance follows the *\_service.\_protocol.name* format as defined in RFC 2782. When you select Free format, enter the entire name in the Domain field.
  - **Service:** Specify the service that the host provides. You can either select a service from the list or type in a service, if it is not on the list. For example, if you are creating a record for a host that provides FTP service, select *\_ftp*. To distinguish the service name labels from the domain name, the service name is prefixed with an underscore. If the name of the service is defined in RFC 1700, Assigned Numbers, use that name. Otherwise, you can use a locally-defined name. This field is disabled when you select **Free Format** as the display input.
  - **Protocol:** Specify the protocol that the host uses. You can either select a protocol from the list or type in a protocol, if it is not on the list. For example, if it uses TCP, select *\_tcp*. To distinguish the protocol name labels from the domain name, the protocol name is prefixed with an underscore. This field is disabled when you select **Free Format** as the display input.
  - **Domain:** If Grid Manager displays a zone name, enter the name here to define an SRV record for a host or subdomain. The displayed zone name can either be the last selected zone or the zone from which you are adding the SRV record. If no zone name is displayed or if you want to specify a different zone, click **Select Zone**. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box. Click a zone name in the dialog box, and then enter the name to define the SRV record. The SRV record name is used to determine the substitute.
  - **Preview:** After you enter all the information, this field displays the FQDN.
  - **DNSView:** Displays the DNS view to which the selected RPZ belongs.
  - **Priority:** Select or enter an integer from 0 to 65535. The priority determines the order in which a client attempts to contact the target host; the domain name host with the lowest number has the highest priority and is queried first. Target hosts with the same priority are attempted in the order defined in the **Weight** field.
  - **Weight:** Select or enter an integer from 0 to 65535. The weight allows you to distribute the load between target hosts. The higher the number, the more that host handles the load (compared to other target hosts). Larger weights give a target host a proportionately higher probability of being selected.
  - **Port:** Specify the appropriate port number for the service running on the target host. You can use standard or nonstandard port numbers, depending on the requirements of your network. You can select a port number from the list or enter an integer from 0 to 65535.
  - **Target:** Enter the canonical domain name of the host (not an alias); for example, *www2.corpxyz.com*.
  - **Comment:** Optionally, enter additional information about the SRV record.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
4. Save the configuration.

### Defining Substitute Rules for TXT Records

A TXT (text) record contains supplemental information for a host. SPF (Sender Policy Framework) records are specialized RPZ TXT records that identify the servers that send mail from a domain.

To define substitute rules for TXT records:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (TXT Record) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Substitute (Record) Rule** -> **Substitute (TXT Record) Rule**.
2. The following fields are displayed in the *Add a Substitute (TXT Record) Rule* wizard:


- **Name:** Enter the name to define a TXT record for a host or subdomain. The name that you specify, irrespective of the RPZ name, is used to determine a match with the RPZ rule. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Text:** Enter the text that you want to associate with the record. It can contain substrings of up to 255 bytes, up to a total of 512 bytes. Additionally, if you enter leading, trailing, or embedded spaces in the text, add quotes around the text to preserve the spaces. For example: " v=spf1 include:corp200.com -all ".
  - **Comment:** Optionally, enter additional information about the TXT record.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
  4. Save the configuration.

### Defining Substitute Rules for IPv4 Addresses or Networks

You can define a substitute for an IPv4 address or a network address. When a client queries for A records of a domain name, if the IP address in A records in the response match the specified address or network, then the response is modified to instead contain the substituted address.

To define substitute rules for IPv4 addresses or networks:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (IPv4 Address) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Substitute (Record) Rule** -> **Substitute (IPv4 Address) Rule**.
2. The following fields are displayed in the *Add a Substitute (IPv4 Address) Rule* wizard:
  - **IP Address or Network:** Enter the IPv4 address which you want to substitute with another IPv4 address. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Policy:** Displays the selected policy.
  - **Substituted IP Address:** Enter the IPv4 address that must be returned to the user when the response matches the A records.
  - **Comment:** Optionally, enter additional information about the IPv4 address.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes. 4. Save the configuration.

 You cannot define a substitute rule for the same IP address or a network address for which you have already defined a passthru rule.

### Defining Substitute Rules for IPv6 Addresses or Networks

You can restrict access to specific IPv6 addresses or networks by providing a substitute IP address. When a client queries for AAAA records of a domain name if the IP addresses in AAAA records in the response match the specified address or network, then the response is modified to instead contain the substituted address.

To define substitute rules for IPv6 addresses or networks:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select *DNS\_View* -> *Zone*, and then click **Add** -> select **Substitute (Record) Rule** -> **Substitute (IPv6 Address) Rule**.  
or  
From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then select a zone. Expand the Toolbar, click **Add** -> select **Substitute (Record) Rule** -> **Substitute (IPv6 Address) Rule**.
2. The following fields are displayed in the *Add a Substitute (IPv6 Address) Rule* wizard:
  - **IP Address or Network:** Enter the IPv6 address or the network address which you want to substitute with another IP address. Click **Select Zone** to select a different zone.
  - **DNS View:** Displays the DNS view to which the selected RPZ belongs.
  - **Policy:** Displays the selected policy.

- **Substituted IP Address:** Enter the IPv6 address that must be returned to the user when the response matches the AAAA records.
  - **Comment:** Optionally, enter additional information about the IPv6 address.
  - **Disable:** Clear the checkbox to enable the record. Select the checkbox to disable it.
3. Click **Next** to define extensible attributes.
  4. Save the configuration.

## Configuring Infoblox Threat Intelligence Feed

Starting with NIOS 7.3.200, Infoblox introduces the Infoblox Threat Intelligence Feed, a threat feed subscription for RPZ updates that offer protection against malicious hostnames. Contact your Infoblox representative for pricing and availability information.

When you upgrade from a previous NIOS release to NIOS 7.3.200 and later releases, the Infoblox RPZ feeds you configured in the previous NIOS release are migrated to the upgraded release. For information about the old RPZ feeds, refer to the NIOS 7.3.4 and earlier *NIOS Administrator Guides*.

You can configure the Threat Intelligence Feed and receive reputation RPZ updates on a regular basis. An RPZ feed receives response policies from the Infoblox in-house threat intelligence team, which produces reputation RPZ data and transfers the data to Grid name servers through zone transfers with or without a TSIG key. To ensure proper authentication and integrity of the RPZ feed zone transfers, using a TSIG key is recommended.



### Note

TSIG Key is used for authentication when downloading information about threat protection feeds. If you have a complex configuration, such as using standalone appliances or Grids that receive threat protection feeds from other standalone appliances or Grids and not directly from the Infoblox distribution servers, ensure that you use the same TSIG key for the RPZ feed zone transfers.

Note that the RPZ feed must have an external primary name server before you can configure it. To propagate RPZs as quickly as possible, the secondary DNS server needs an address to which the RPZ source feed can send NOTIFY messages. For example, if the secondary DNS server is configured behind a NAT, you may want to establish a one-to-one NAT for the lead secondary DNS server so it can receive NOTIFY messages from the RPZ source feed. Otherwise, the lead secondary DNS server will need to periodically poll the RPZ source feed, which could take longer than expected.



### Note

To enter IDNs (Internationalized Domain Name) in an RPZ feed, you can use the punycode representation of the IDN.

To configure the Threat Intelligence Feed:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then click the *Add* icon.
2. When you click the *Add* icon, either the *Add Response Policy Zone Wizard* or the *Add DNS View* wizard is displayed based on the following:
  - When you click the *Add* icon, the *Add Response Policy Zone Wizard* is displayed, if you have not created additional *DNS views* and only have the *default view*.
  - If you have configured multiple DNS views, you must drill-down to the corresponding *DNS\_View* to assign an RPZ feed. Click the *Add* icon and the *Add Response Policy Zone Wizard* is displayed. To create a new DNS view for your RPZ feed, click the *Add* icon and complete the details in the *Add DNS View* wizard. For information about adding and modifying a DNS View, see [Configuration Example: Configuring a DNS View](#).
3. In the *Add Response Policy Zone Wizard*, select **Add Response Policy Zone Feed**, click **Next** and specify the following:
  - **Name:** Enter the name of the Infoblox RPZ feed. It can be a combination of alphanumeric characters. You can enter up to 256 characters. For more information, see Infoblox Threat Intelligence Feeds below.
  - **DNS View:** The name of the view that you have selected is displayed by default. You can select a view from the drop-down list to associate it with the RPZ feed.

- **Policy Override:** Select a value from the drop-down list. You can override the policy actions that are specified in the rule level.
- **Log Only (Disabled)** – Select this if you want to disable an RPZ rewrite using rules in the RPZ zone. If the response to the recursive query matches any RPZ rule, the rule is logged, but the response will not be altered. You cannot overwrite the response to the user. Note that this option will not override RPZ rules in other RPZ zones, if they take precedence.

Note that when you select the **Log Only** option, the RPZ related reports are not updated even though the information is logged to the syslog.

- **None (Given)** – Select this if you want to use the policy from the rule level.
  - **Block (No Data)** – Select this if you want the user to receive a response that indicates that there is no data.
  - **Block (No Such Domain)** – Select this if you want the user to receive a NXDOMAIN as the DNS response. All the policy actions in an RPZ are replaced with a NXDOMAIN block.
  - **Passthru** – Select this if you want the user to see the actual response without modification. All the policy actions in an RPZ are replaced with the passthru action.
  - **Substitute (Domain Name)** – Select this if you want to replace all the policy actions in an RPZ with the substitution action that is specified.
    - **Domain Name:** This appears only when you select **Substitute (Domain Name)** from the **Policy Override** list. Enter the domain name that you want the client to receive instead of the actual domain name, which is malicious or unauthorized.
  - **Severity:** Select the threat severity level for the RPZ zone. The threat severity you select here determines the severity for the RPZ rule. Select **Critical**, **Major**, **Warning**, or **Informational**. The default threat severity level is Major. Note that each of these levels is represented by a number in the syslog (8 being Critical and 4 being Informational). When you upgrade to NIOS 7.0.0, the appliance automatically updates the threat severity level to Informational (displayed as 4 in the syslog) for existing RPZ zones. For information about viewing RPZ in the syslog and severity levels, see [Verifying RPZ Configuration](#).
  - **Comment:** Optionally, enter additional information about the Infoblox RPZ feed.
  - **Disable:** Select the checkbox to disable the RPZ feed without deleting its configuration. Clear the checkbox to enable the RPZ feed. For information, see [Enabling and Disabling Zones](#). Note that disabling an RPZ feed may take a longer time to complete depending on the size of the data.
  - **Lock:** Select the checkbox to lock the RPZ feed so that you can make changes to it and prevent others from making conflicting changes. For information about Locking and Unlocking RPZs, see [Managing RPZs](#).
4. Click **Next** to associate the RPZ feed with at least one external primary name server and a secondary name server:
- Define name servers for the RPZ feed. An RPZ feed must have at least one RPZ source as an external primary name server and at least one Grid secondary name server. For external primary servers, specify the following:
    - **Name:** Enter the zone name of the primary name server.
    - **Address:** Enter the name server IP address provided by Infoblox for the RPZ feed.
    - **Use TSIG:** Select the checkbox to specify TSIG settings.
    - **Key Name:** Enter the TSIG Key Name provided by Infoblox.
    - **Key Algorithm:** Select **hmac-md5**.
    - **Key Data:** Enter the TSIG string provided by Infoblox.

Note that either the Grid name server or the DNS view must be recursive for the RPZ feed. You can associate a lead secondary with an RPZ feed. For information on specifying primary and secondary, see [Assigning Zone Authority to Name Servers](#). When you select **All Recursive Name Servers** from the list, all the recursive name servers in the Grid are added as secondary servers for the zone. For information about how to configure a local RPZ, or RPZ feed, or FireEye RPZ for all recursive servers, see [Configuring RPZs for All Recursive Servers](#). For information on specifying name server groups, see [Using Name Server Groups](#).
5. Save the configuration and click **Next** to define extensible attributes. Click **Restart** if it appears at the top of the screen. For information, see [Managing Extensible Attributes](#).



## Infoblox Threat Intelligence Feeds

Infoblox RPZ feeds are categorized into pure malicious feeds and combination feeds. All the feeds listed below are set to return NXDOMAIN for items in the feed. Threat data changes are pushed every 20 minutes from the DNS servers and significant changes are typically made every two hours.

The following table list the Infoblox Threat Intelligence feeds:

### *Pure Malicious Feeds*

Name	Description
Base (base.rpz.infoblox.local)	Enables protection against known hostnames that are dangerous as destinations, such as APT, Bot, Compromised Host/Domains, Exploit Kits, Malicious Name Servers, and Sinkholes along with bogon IP addresses.
AntiMalware (antimalware.rpz.infoblox.local)	Enables protection against known malicious threats that can take action on or control of your system, such as Malware Command & Control, Malware Download, and active Phishing sites.
Ransomware (ransomware.rpz.infoblox.local)	Enables protection against ransomware that restricts access to the computer system that it infects, and demands a ransom paid to the creator of the malware in order for the restriction to be removed. Some forms of ransomware encrypt files on the system's hard drive, while some may simply lock the system and display messages intended to coax the user into paying. Examples include Locky, CryptoLocker, Dircrypt, and CryptoWall.
Bogon (bogon.rpz.infoblox.local)	Enables protection against bogons, which are commonly found as the source addresses of DDoS attacks. A bogon is an informal name for an IP packet on the public Internet that claims to be from an area of the IP address space reserved, but not yet allocated or delegated by the Internet Assigned Numbers Authority (IANA) or a delegated Regional Internet Registry (RIR). The areas of unallocated address space are called bogon space. Many ISPs and end-user firewalls filter and block bogons, because they have no legitimate use, and usually are the result of accidental or malicious misconfiguration.
AntiMalware_IP (antimalware-ip.rpz.infoblox.local)	Enables protection against known malicious or compromised IP addresses. These are known to host threats that can take action on or control of your system, such as Malware Command & Control, Malware Download, and active Phishing sites.
Bot_IP (bot-ip.rpz.infoblox.local)	Enables protection against self-propagating malware designed to infect a host and connect back to a central server or servers that act as a command and control (C&C) center for an entire network of compromised devices, or "botnet." With a botnet, attackers can launch broad-based, "remote-control," flood-type attacks against their target(s). Bots can also log keystrokes, gather passwords, capture and analyze packets, gather financial information, launch DoS attacks, relay spam, and open back doors on the infected host.
ExploitKit_IP (exploitkit-ip.rpz.infoblox.local)	Enables protection against distributable packs that contains malicious programs that are used to execute "drive-by download" attacks in order to infect users with malware. These exploit kits target vulnerabilities in the users' machines (usually due to unpatched versions of Java, Adobe Reader, Adobe Flash, Internet Explorer, ...) to load malware onto the victim's computer.
Malware_DGA (malware-dga.rpz.infoblox.local)	Domain generation algorithm (DGA) are algorithms seen in various families of malware that are used to periodically generate a large number of domain names that can be used as rendezvous points with their command and control servers. Examples include Ramnit, Conficker, and Banjori.



Name	Description
TOR_Exit_Node_IP (tor-exit-node-ip.rpz.infoblox.local)	Tor Exit Nodes are the gateways where encrypted Tor traffic hits the Internet. This means an exit node can be used to monitor Tor traffic (after it leaves the onion network). It is in the design of the Tor network that locating the source of that traffic through the network should be difficult to determine
SURBL_Multi (multi-domain.surbl.rpz.infoblox.local)	Blacklist of Malicious Domains including up-to-date intel on active malware, phishing, botnet, and spam domains. Based on data provided by our partner SURBL.
SURBL_Fresh (fresh-domain.surbl.rpz.infoblox.local)	Newly Observed Domains. SURBL Fresh feed provides critical, accurate, information on the time new domains are placed into service. Security policy can be easily applied (block, quarantine, walled garden, etc.) to prevent resolution of new domains, based on the user's defined policies. Based on data provided by our partner SURBL.

## Downloading Rules for an RPZ Feed

You can perform a zone transfer to transfer the rules from an external primary name server to the RPZ feed. You cannot modify these rules, but you can override the entire ruleset or an individual rule. However, if you import a zone to a local zone, you can edit the rules within a local zone. The feed zone supports NSIP and NSDNAME rules; however local RPZs do not support these rules. To download rules from an external primary name server:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then click the corresponding *RPZ Feed*.
2. In the **Export** dialog box, complete the following:
  - **Separator**: Select the separator used in the data file. The default value is **Comma**.
  - Click **Export**.

All the rules are transferred. You can download rules only if the lead secondary has completed at least one zone transfer from the external primary. You can either open the data file or save it to your computer. The rules are displayed for the selected RPZ feed in the *Rule* wizard.

After you have downloaded rules from an RPZ feed, you can test RPZ feed policies, as described in [Testing RPZ Feed Rules](#).

## Testing RPZ Feed Rules

After you have downloaded rules from an RPZ feed, you can test the downloaded policies by using the dig command and observing log messages that contain redirect or rewrite responses in the syslog. The NIOS appliance supports generation of RPZ log messages in CEF (Common Event Format). Note that non-RPZ messages cannot be generated in CEF.

You must enable the **rpz** option in the **Logging Category** of the *Grid DNS Properties* editor to receive RPZ related messages in the syslog. For information about setting DNS logging categories, see [Using a Syslog Server](#).

To view RPZ log messages in the syslog, you can use the system filter **RPZ Logs** from the **Quick Filter** to filter the messages. Note that only messages in CEF are displayed.

To view RPZ log messages:

1. From the **Administration** tab, select the **Logs** tab -> **Syslog** tab.
2. From the drop-down list at the upper right corner, select the Grid member on which you want to view the syslog.
3. Click **Show Filters** to enable the filters. Select **RPZ Logs** from the **Quick Filter** drop-down list to narrow down the system messages you want to view.

The name server recursive cache makes a syslog entry when an RPZ functionality fails. The syslog message log format is as follows:

```
rpz <TYPE> rewrite <QUERY> via <RPZ_RECORD><ERROR_MESSAGE>
```

where: <TYPE> is one of following RPZ action types: QNAME, IP, NSIP, NSDNAME, CLIENT-IP;

<QUERY> is a query record to process;

<RPZ\_RECORD> is an RPZ record that is used to perform an action to the query;

<ERROR\_MESSAGE> is a message with error details. Example: NS address rewrite rreset failed:, concatenate() failed:, NS db\_find() failed:, stop on qresult in rpz\_rewrite() failed:, stop on unrecognized qresult in rpz\_rewrite() failed:, etc.

To test RPZ feed policies:

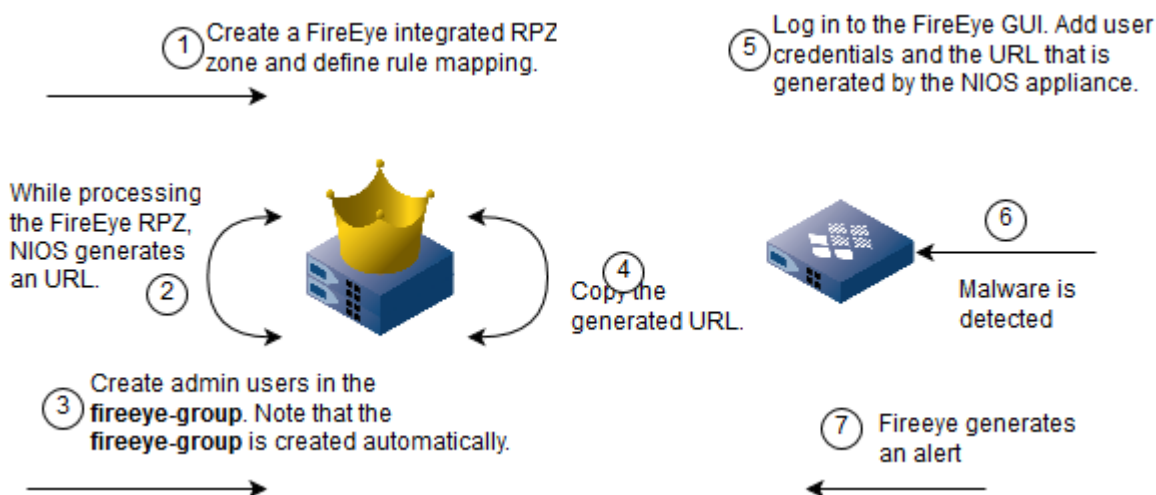
1. Open a terminal console on your computer.
2. Type the command **dig @<your DNS server IP> <queried domain>**.
3. Go to the **Administration** tab -> **Logs** tab -> **Syslog** tab to view CEF log messages.

## About FireEye Integrated RPZs

Infoblox DNS Firewall provides a mechanism to further protect your network from malware and APTs (Advanced Persistent Threats) through the integration of FireEye appliances. When your NIOS appliance is properly integrated with a FireEye appliance, it receives periodic alerts and APTs from the FireEye appliance when it identifies such threats. Based on your configuration, the NIOS appliance translates these alerts into RPZ rules that not only further protect your network from malicious attacks, but also aid in identifying clients that have been compromised.

As illustrated in the below figure, after installing the required RPZ and Security Ecosystem licenses on the NIOS appliance, you can configure a FireEye integrated RPZ in which you map RPZ rules to FireEye alert types. While creating the FireEye RPZ, the appliance generates a URL to which the FireEye appliance sends alerts. Ensure that you enter this URL when configuring the FireEye appliance. The NIOS appliance also creates the **fireeye-group** admin group after you define the first FireEye RPZ. You can add multiple admin users to this admin group. Note that users in the **fireeye-group** can only send alerts to the NIOS appliance; they cannot access the Infoblox GUI, CLI, API and RESTful API. They also do not have permissions to perform other tasks on the appliance. Ensure that you record the usernames and passwords for all user accounts so you can enter them correctly when you configure the FireEye appliance. You can map a single or multiple FireEye appliances to a NIOS appliance where multiple users or zones exist.

*FireEye Integrated RPZ*



To configure a FireEye integrated RPZ, complete the following:

1. Create a new FireEye integrated RPZ, as described in Configuring FireEye RPZs below.
2. Create FireEye admin users, as described in For FireEye Integrated RPZs below.

3. Add URL and user credentials on the FireEye appliance, as described in [Configuring the FireEye appliance](#) below.
4. When a malware or threat is detected, the FireEye appliance sends an alert message to the NIOS appliance, which is stored in the syslog. For more information, see [Handling Alerts from the FireEye appliance](#) below.

## Configuring FireEye RPZs

You must create an RPZ zone and map the FireEye alerts with an RPZ rule to receive alerts from FireEye. These alerts will then be translated into appropriate RPZ rules that are added to the FireEye RPZ. You can also define a time limit for a specific alert type or set the alert type to live forever. When you define a lifetime, the alert type will be active for the specified number of days or weeks in the NIOS appliance, and will then expire after the specified time. After you configure the FireEye integrated RPZ, the NIOS Grid receives alerts from the FireEye appliance and creates RPZ rules for some of the alerts received. FireEye appliance sends alert messages with basic authentication. You must configure a username and password on the NIOS appliance prior to receiving any alerts from the FireEye appliance.



### Note

The NIOS appliance treats the FireEye integrated RPZ as a local RPZ. Thus, you cannot assign an external primary name server to the zone.

An alert contains the malware URL along with a valid FQDN. The NIOS appliance can only map an alert to a RPZ rule if the FQDN is present. If an alert doesn't contain the FQDN, then the alert is ignored by the NIOS appliance. Once the alert is processed and properly mapped to an RPZ rule, it remains in the database until you delete it manually. You can get more information about the alerts, which are sent by the FireEye appliance, from the syslog.



### Note

You can configure feeds from multiple FireEye appliances. To enable or disable FireEye integration module feeds from individual appliances, you must enable or disable user access of the particular FireEye appliance. Note that the FireEye feeds will not be in the RPZ format, but when you configure a FireEye integrated RPZ, the NIOS appliance creates a new URL through which the FireEye appliance sends alerts.

To configure a FireEye integrated RPZ:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, and then click the *Add* icon.
2. When you click the *Add* icon, either the *Add Response Policy Zone Wizard* or the *Add DNS View* wizard is displayed based on the following:
  - When you click the *Add* icon, the *Add Response Policy Zone Wizard* is displayed if you have not created additional DNS views and only have the default view.
3. If you have configured multiple DNS views, you must drill-down to the corresponding view to assign a FireEye Integrated RPZ. Click the *Add* icon and the *Add Response Policy Zone Wizard* is displayed. To create a new DNS view for your FireEye integrated RPZ, click the *Add* icon and complete the details in the *Add DNS View* wizard. For information adding and modifying a DNS view, see [Adding a DNS view](#).
4. In the *Add Response Policy Zone Wizard*, select **Add FireEye-Integrated Response Policy Zone**, click **Next** and specify the following:
  - **Name:** Enter the name of the FireEye integrated RPZ. It can be a combination of alphanumeric characters. You can enter up to 256 characters.
  - **DNS View:** The name of the view that you have selected is displayed by default. You can select a view from the drop-down list to associate it with the FireEye integrated RPZ.
  - **Policy Override:** Select a value from the drop-down list. You can override the policy actions that are specified in the rule level.
    - **Log Only (Disabled)** – Select this if you want to disable an RPZ rewrite using rules in the RPZ. If the response to the recursive query matches any RPZ rule, then the rule is logged, but the response will not be altered. Note that this option will not override RPZ rules in other RPZ zones, if they take precedence. Select this option to preview the rules in the syslog before they take effect.
    - **None (Given)** – Select this if you want to use the policy from the rule level.

- **Block (No Data)** – Select this to send a response that contains no data in it.
- **Block (No Such Domain)** – Select this if you want the user to receive a DNS response that indicates there is no domain. All the policy actions in an RPZ are replaced with a NXDOMAIN block.
- **Passthru** – Select this if you want to send an actual response without modification. All the policy actions in an RPZ are replaced with the passthru action.
- **Substitute (Domain Name)** – Select this if you want to replace all the policy actions in an FireEye integrated RPZ with the specified substitution action.
  - **Domain Name:** This appears only when you select **Substitute (Domain Name)** from the **Policy Override** list. Enter the domain name that you want the client to receive instead of the actual domain name, which is malicious or unauthorized.
- **Comment:** Optionally, enter additional information about the FireEye integrated RPZ.
- **Disable:** Select the checkbox to disable the FireEye integrated RPZ without deleting its configuration. Clear the checkbox to enable the FireEye integrated RPZ. For information, see [Enabling and Disabling Zones](#). Note that disabling the FireEye integrated RPZ may take a longer time to complete depending on the size of the data.

5. Click **Next** to define rule mapping:

- **Server URL:** The appliance displays the URL that you use when configuring the FireEye appliance. This URL is used to handle alerts, which is sent by the FireEye appliance. It handles alerts based on the standard authentication. The URL generated by the NIOS appliance consists of the Grid Manager IP address, network view, and DNS view of the FireEye zone. If you change the IP address, network view, zone or DNS view after you have configured a FireEye RPZ, the URL will change accordingly. Thus FireEye will not be able to send alerts to the updated URL. You must update the URL in the FireEye appliance to send alerts to the NIOS appliance. The **Server URL** is generated in this format:

`https://<host address>/alert/feye/<network view>/<dns view>/<zone>`

- **Rule Mapping:** You can map a **FireEye alert type** with an RPZ policy. Select an **RPZ policy type** from the drop-down list. Note that the **FireEye alert type** is read-only. The NIOS appliance applies corresponding RPZ policy type when the FireEye appliance sends an alert to the NIOS appliance. You can also specify a time limit for each FireEye RPZ rule depending on the FireEye alert type. NIOS displays default lifetime value for each alert type. You can change the default lifetime of the alert type. When you define a value, the value must be greater than zero. When you select **Live Forever** from the drop-down list, the alert type will never expire and will be stored in the database until further notice. The NIOS appliance will use the default time if you do not specify a value. You can specify the expiration time in days or weeks only. The following table lists the FireEye alerts, RPZ policy types, and the time limit for a specific FireEye alert:

#### FireEye Rule Mapping

FireEye Alert Type	RPZ Policy Type	Lifetime
Domain Match Infection Events Callback Events Malware Object Web Infection	Select a value from the drop-down list for a FireEye alert when a malware object is detected: <b>None</b> , <b>Passthru</b> , <b>Block (No Such Domain)</b> , <b>Block (No Data)</b> , and <b>Substitute (Domain Name)</b> . The drop-down list displays <b>Passthru</b> , by default. For more information about the <b>RPZ Policy Types</b> , see <a href="#">Configuring Rules for RPZs</a> .	Specify a lifetime for each FireEye alert in <b>Days</b> , <b>Weeks</b> , or select <b>Live Forever</b> from the drop-down list. The following are the default values for different alert types: <ul style="list-style-type: none"> <li>• Domain Match - 1 week</li> <li>• Infection Events - 1 day</li> <li>• Callback Events - 1 week</li> <li>• Malware Object - 1 day</li> <li>• Web Infection - 1 day</li> </ul> Click on the default value to change the lifetime value.

When you edit the lifetime of an existing alert type, NIOS deletes the alert type based on the new lifetime setting. It also updates the expiration time for the corresponding alert type. Note that there might be an impact on the performance when you delete expired FireEye RPZ rules.

- **Override rule mapping for APT events:** Select a value from the drop-down list to override rule mapping for Advanced Persistent Threats. Events that are marked as APT events by FireEye override rules that are set for other event types. The values in the drop-down list are:

- **NoOverride** – Select this if you want to use the policy from the rule level and do not want to override the rule mapping settings. This value is displayed in the drop-down list, by default.
  - **Passthru** – Select this if you want the user to see the actual response without modification. All the policy actions in an RPZ are replaced with the passthru action.
  - **Block (No Such Domain)** – Select this if you want the user to receive a NXDOMAIN as the DNS response. All the policy actions in an RPZ are replaced with a NXDOMAIN block.
  - **Block (No Data)** – Select this if you want the user to receive a response that indicates that there is no data.
    - **Substitute (Domain Name)** – Select this if you want to replace all the policy actions in an RPZ with the substitution action that is specified.
  - **Substituted Domain Name:** This appears only when you select **Substitute (Domain Name)** from the **Policy Override** list either for APT events or for FireEye alerts. Enter the domain name that you want the client to receive instead of the actual domain name, which is malicious or unauthorized.
6. Click **Next** to associate the FireEye integrated RPZ with at least one primary name server:
- Define the name servers for the FireEye integrated RPZ. A Grid name server must be recursive when primary Grid name server is used as an RPZ source. A FireEye integrated RPZ may or may not have a recursive server. For example, there could be a Grid that has only primary Grid name server for a FireEye integrated RPZ to act as an RPZ source for an external set of name servers. When you select **All Recursive Name Servers** from the list, all the recursive name servers in the Grid are added as secondary servers for the zone. For information on specifying primary or secondary name servers, see [Assigning Zone Authority to Name Servers](#). For information on specifying name server groups, see [Using Name Server Groups](#). For information about all recursive name servers, see [Configuring RPZs for All Recursive Servers](#).
7. Save the configuration and click **Next** to define extensible attributes.
8. Click **Restart** if it appears at the top of the screen.

## Configuring Rules for FireEye RPZs

You can define a list of rules based on how the DNS server determines its response to recursive queries. Based on the rules defined, responses to clients are either manipulated or forwarded without any changes. To configure rules for FireEye RPZs:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, click *DNS\_View* -> *Zone* and then click **Add** -> select a **Rule**.
2. The rules are classified as follows, for more information about these rules see, [Configuring Rules for RPZs](#):
  - **Passthru Rule**
  - **Block (No Such Domain) Rule**
  - **Block (No Data) Rule**
  - **Substitute (Domain Name) Rule**
  - **Substitute (Record) Rule**
3. Complete the details in the corresponding editor.
4. Save the configuration and click **Next** to define extensible attributes.

## Configuring the FireEye appliance

You must configure the FireEye appliance to send alerts to the NIOS appliance. Ensure that the following are complete before you configure the FireEye appliance:

1. Install required license on the NIOS appliance. For more information about license, see [License Requirements and Admin Permissions](#).
2. Create a new FireEye RPZ zone.
3. Create FireEye admin users. For more information, see [For FireEye Integrated RPZs](#).
4. Get the URL from the NIOS appliance and record it. You need this to configure the FireEye appliance. For more information about the Server URL, see [Configuring FireEye RPZs](#) above. If you have already configured a FireEye integrated RPZ, then you can retrieve the URL through the **FireEye** tab of the corresponding FireEye RPZ zone. For more information about managing and retrieving the URL, see [Modifying RPZs](#).

5. Record the usernames and passwords on the NIOS appliance. You must use these credentials when configuring FireEye alerts to enable the alerts to be received by NIOS. For more information, see [Configuring the FireEye appliance to send alerts to NIOS](#) below.

### Configuring the FireEye appliance to send alerts to NIOS

You must configure the NIOS generated URL, usernames and passwords on the FireEye appliance. FireEye appliance embeds the configured usernames and passwords in the alerts for authentication. When an alert is received, the NIOS appliance verifies the FireEye username prior to processing the alert. Note that the NIOS appliance accepts alerts sent by the FireEye appliance in **JSON Normal** format only.

To configure a FireEye appliance:

1. Login to the FireEye appliance with your username and password.
2. In the FireEye GUI, click **Settings** tab and then click the **Notifications** tab on the left panel.
3. In the **Notification Settings** page, click the **http** link and then enter the name of the HTTP server you want to add. Click **Add HTTP Server** and complete the following:
  - **Name:** When you click add, the HTTP server name that you specified is listed in this column.
  - **Enabled:** Select the checkbox to enable alerts and notifications for the HTTP server.
  - **Server Url:** Enter the URL you received on the NIOS appliance. The alerts and notifications are sent using this URL by the FireEye appliance.
  - **Auth:** Select this checkbox if authentication is required for the server.
  - **Username and Password:** Enter the Username and Password of the user that you have configured for the **fireeye-group** on the NIOS appliance. For more information, see [For FireEye Integrated RPZs](#).
  - **Notification:** Select a notification from the drop-down list. You can choose to include notifications for all events or only events of a selected type. The FireEye appliance will send an alert to the NIOS appliance only when selected event is encountered. When you select **All Events**, alerts are sent when each event is encountered by the FireEye appliance.
  - **Delivery:** Select **Per Event** from the drop-down list. Note that the NIOS appliance supports only **Per Event** selection. The FireEye appliance sends an alert each time it encounters an event.
  - **Account:** You can specify a user account name for this notification.
  - **SSL Enable:** Select this checkbox to enable SSL for secure transmission of alerts from the FireEye appliance to NIOS.
  - **Default Provider:** Select a default provider from the list.
  - **Message Format:** Select **JSON Normal** from the drop-down list. Note that the NIOS appliance supports only this message format.
4. Click **Update** at the bottom of the page.

#### **Note**

You can also click **Test-Fire** to test the configuration. If the configuration is successful, FireEye sends a confirmation message to the NIOS appliance and the NIOS appliance logs this message in the syslog. It generally takes a few seconds for the NIOS appliance to receive alerts. You must verify the configuration, if there is no entry in the syslog.

### Handling Alerts from the FireEye appliance

The NIOS appliance processes each alert that it receives from the FireEye appliance. The alert contains the malware URL along with a valid FQDN. NIOS appliance can only map an alert to a RPZ rule if the FQDN is present. Once the alert is processed and properly mapped to an RPZ rule, it remains in the database until you delete it manually. When the RPZ rule is different from the existing rules, the new RPZ rule gains precedence over the existing RPZ rule in the FireEye integrated RPZ. Note that you cannot retrieve alerts that are ignored. You can get more information about the alerts, which are sent by the FireEye appliance, from the syslog. An alert will not be processed and will be ignored:

- when there are changes to the URL or if the alert does not have the malware URL or FQDN in them.
- if the zone is not found.
- if the alert is sent without any username in it or if the username does not belong to the fireeye-group.



- if a FireEye admin user is deleted. NIOS will neither authenticate the deleted user credentials nor process any future alerts with deleted user credentials.
- if the search mapping fields contain IP addresses other than FQDNs.
- if alerts contain domain names in an IPv4 or IPv6 address format.

## Logging FireEye Integrated RPZ messages

The NIOS appliance logs FireEye events and alerts in the syslog and audit log. Each FireEye feed event is logged every time an alert is sent to NIOS by the FireEye appliance. When you create a new rule or update an existing rule, then those are also logged in the syslog. You can use messages logged in syslog to verify events that are related to communication between the FireEye and NIOS appliances. It also enables the admin to monitor alerts and verify how the alerts are processed. Details about alerts that are received and processed are also logged. Syslog messages are logged when:

- an alert is received from the FireEye appliance.
- syslog messages contain required information for reporting.
- an alert is successfully mapped to an RPZ rule. The message format is as follows:
  - <FireEye: Found an APT alert>
- the NIOS appliance cannot process alerts. For example, alert structure mismatch, unrecognizable data, etc. The messages will have the following format:
  - <FireEye: Cannot parse FQDN due to missing field"cnc-services">
  - <FireEye: Cannot determine if it is an APT alert..>
  - <FireEye: Invalid Alert Type ....>
  - <FireEye: Couldn't find the required field...>
  - <FireEye: No mapping rule has been set for alert type.....>
- a duplicate alert is sent by the FireEye appliance for which the same RPZ rule already exists.

### Note

For debugging purposes, alert messages will be displayed in the infoblox.log file.

NIOS periodically scans the syslog of a member that has RPZ license installed to generate recent hits data for the *RPZ Recent Hits* tab. This might cause a performance impact as CPU cycles will be used on the member. For more information about *RPZ Recent Hits* tab, see [RPZ Recent Hits](#).

## Configuration Examples

This section illustrates some of the examples of local and FireEye integrated RPZs.

### Local RPZ Examples

Following is an example of an IP related rule. For example, execute the following command:

```
dig @10.35.104.19 abc.net
```

If the above command returns `18.58.20.1`, then define an IPv4 substitute rule `18.0.0.0/8`, Substitute (IPv4) `8.8.8.89`.

Execute the command `dig @10.35.104.19 abc.net` again. You will receive the substituted address instead of the actual domain name.

Following is an example of values in CEF for the above substitution example:

```
2012-11-06T19:04:02+00:00 daemon (none) named[25193]: info
```

```
CEF:0|Infoblox|NIOX|6.6.0-185622|RPZ-IP|records|4|app=DNS dst=10.35.104.19
src=10.32.0.242 spt=50035 view=_default qtype=A msg="rpz IP records rewrite
abc.net [A]
via 8.0.0.0.18.rpz-ip.localrpz"
```

Following is an example of values in the CEF for Block (No Data) rules:

```
2012-11-06T19:00:01+00:00 daemon (none) named[25193]: info
CEF:0|Infoblox|NIOX|6.6.0-185622|RPZ-QNAME|NODATA|4|app=DNS dst=10.35.104.19
src=10.32.0.242 spt=50035 view=_default qtype=A msg="rpz QNAME NODATA
rewrite nodata.net [A]
via nodata.net.localrpz"
```

You can view the NIOS version, name of the view, source, and destination.

#### FireEye Integrated RPZ Examples

Following is an example of a syslog message when an alert gets converted to an RPZ rule:

```
013-09-11T10:59:55-07:00 user (none) httpd[]: info fireeye-rpt:
'79167','infection-match','minr''eng-lab-249.inca.infoblox.com'
2013-09-11T10:59:55-07:00 user (none) httpd[]: info FireEye: Create an RPZ
rule for
'd.bnksw.com' with 'SUBSTITUTE' rule in RPZ zone 'com.lock'
```

Note that the domain name lock.com is displayed in the reverse format.

Following is an example of a syslog message when an alert is ignored by the NIOS appliance:

```
2013-09-11T11:04:01-07:00 user (none) httpd[]: info fireeye-rpt:
'114488','malware-object','majr''eng-lab-249.inca.infoblox.com'
2013-09-11T11:04:01-07:00 user (none) httpd[]: info FireEye: Cannot parse
FQDN due to
missing field''cnc-services''
```

#### Example of a basic RPZ Workflow

Following is an example of a basic RPZ workflow:

1. Install the RPZ license. For more information, see [License Requirements and Admin Permissions](#).
2. Enable recursive queries for a DNS view, member, or Grid, as described in [Enabling Recursion for RPZs](#).
3. Enable RPZ logging in the *Grid DNS Properties* editor to view syslog entries for RPZ queries. For more information about setting DNS logging categories, see [Using a Syslog Server](#).



4. Create a local RPZ. For more information, see [Configuring Local RPZs](#).
5. Define a **Substitute (PTR Record) Rule** for domain name 3.3.3.5.in-addr.arpa, which is substituted with the domain name ptr1.com. For more information, see [Defining Substitute Rules for PTR Records](#).
6. Execute the dig command to view output. The output contains the substituted domain name ptr1.com. Following is the output of an RPZ query for **Substitute (PTR Record) Rule**:

```
$ dig @10.36.2.73 3.3.3.5.in-addr.arpa in ptr
; <<>> DiG 9.6.2-P2-RedHat-9.6.2-5.P2.fc12 <<>> @10.36.2.73 3.3.3.5.in-addr.arpa in ptr
; (1 server found)
;; global options: +cmd
;; Got answer:
;; -->HEADER<<- opcode: QUERY, status: NOERROR, id: 7351
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 0, ADDITIONAL: 0

;; QUESTION SECTION:
;3.3.3.5.in-addr.arpa.          IN      PTR

;; ANSWER SECTION:
3.3.3.5.in-addr.arpa. 7200    IN      PTR      ptr1.com.

;; Query time: 3 msec
;; SERVER: 10.36.2.73#53(10.36.2.73)
;; WHEN: Thu Sep 26 23:27:10 2013
;; MSG SIZE rcvd: 60
```

7. Following is the syslog entry for the query mentioned above:

```
2013-09-27T02:26:46-04:00 daemon (none) named[21737]: info
CEF:0|Infoblox|NIOs|6.9.0-218052|RPZ-QNAME|Local-Data|4|app=DNS
dst=10.36.2.73
src=10.120.20.194 spt=40518 view=2 qtype=PTR msg="rpz QNAME Local-Data
rewrite
3.3.3.5.in-addr.arpa [PTR] via 3.3.3.5.in-addr.arpa.local1.com"
```

For more information about syslog, see [Viewing RPZ in the Syslog](#).

## Mitigating Cyber Threats using TAXII

To mitigate the increasingly complex cyber attacks, you can enable the appliance to run a TAXII (Trusted Automated eXchange of Indicator Information) service to receive information on real-time threat incidents. The information in each threat incident is represented using the STIX (Structured Threat Information eXpression) language format. STIX is a standard language used to describe structured cyber threat information, which is shared between different TAXII clients. When you run the TAXII service on a Grid member, the appliance acts as a TAXII server that receives TAXII messages

(for one or more specified STIX collection) from TAXII clients. The TAXII message typically contains a list of IP addresses (both IPv4 and IPv6) and domains. The member then communicates with the Grid Master and sends a request to create an RPZ rule on the specified RPZ based on the TAXII messages it receives. The RPZ rule created on NIOS is available in the **Response Policy Zones** tab, as shown in the RPZ Rules created for the Mapped RPZ and Collection figure below.

 **Note**

Once you start the TAXII server, the inbox for the configured collections is available at <https://<member address>/services/inbox> and the TAXII discovery service is available at <https://<member address>/services/discovery>, where <member address> is the MGMT or LAN IP address (IPv4 or IPv6 address of the port that is configured).

For more information about TAXII and STIX, refer to the following:

<https://taxii.mitre.org/>

<http://taxiiproject.github.io/>

### Supported Appliances for TAXII Service

You can run the TAXII service on the following Infoblox appliance models: IB-1410, IB-1415, IB-1420, IB-1425, IB-VM-1410, IB-VM-1415, IB-VM-1420, IB-VM-1425, TE-815, TE-2215, TE-2225, IB-VM-4010, IB-4030-10GE, IB-VM-2220, IB-VM-2225, PT-1405, PT-2205 and PT-2205-10GE.

### Licensing Requirements and Permissions

To enable the TAXII service, you must install the **Security Ecosystem** license on any Grid member. You must also install an **RPZ** license on any Grid member in the Grid in order to create RPZ rules based on the TAXII messages. To allow a group to access the TAXII service, you can enable the group to authenticate with the TAXII server.

To enable a group to access the TAXII server:

1. From the **Administration** tab, select the **Administrators** tab -> **Groups** tab, and then click the Add icon.
2. In the *Add Admin Group* wizard, click the **Roles** tab and then complete the following in the **Allowed Interfaces** section:
  - **TAXII**: Select this checkbox to enable a group to authenticate with the TAXII server.
3. Save the configuration.

### Mapping RPZs with TAXII Collections

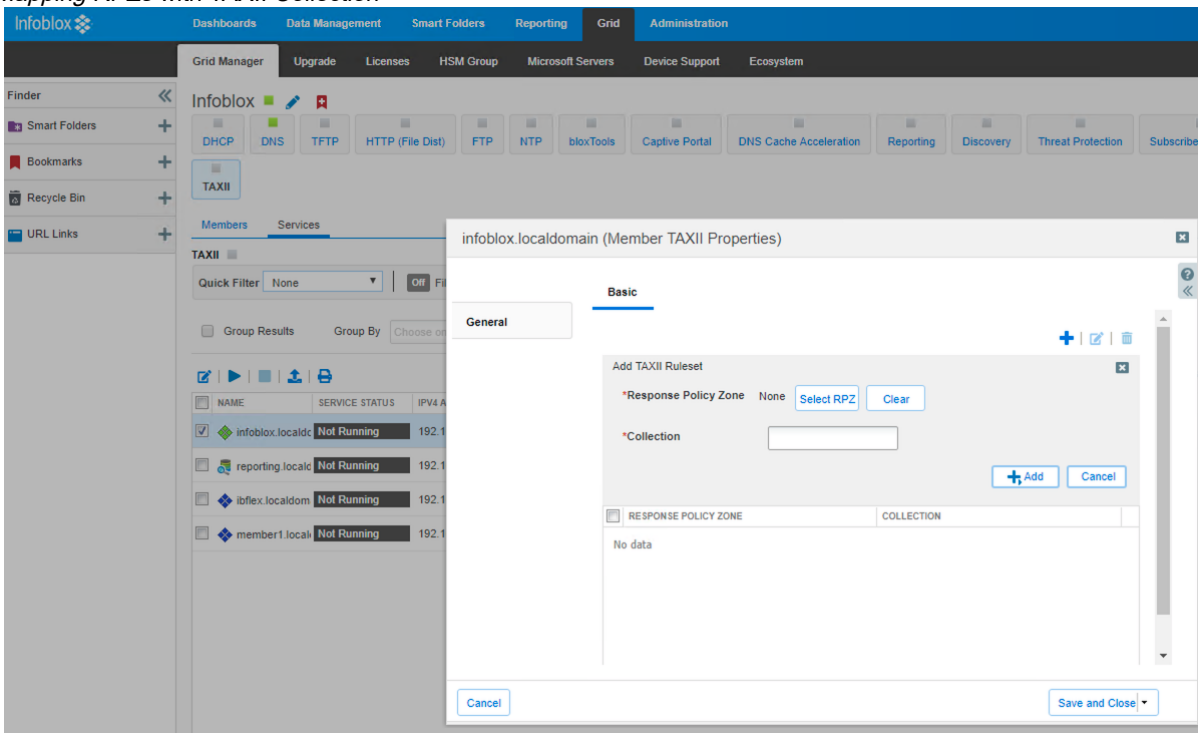
NIOS supports arbitrary set of RPZ rules mapped to the corresponding TAXII collection. To map an RPZ with a TAXII collection:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Services** tab.
2. In the **Services** tab, select the *TAXII\_member* checkbox, and then click **Edit** -> **Member TAXII Properties** from the Toolbar.
3. In the *Member TAXII Properties* editor, complete the following:
  - **Response Policy Zone**: Click the Add icon and click **Select RPZ** to select an RPZ. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box from which you can select one.
  - **Collection**: Enter the name of the TAXII collection that will be mapped to the RPZ. Note that you can only use valid URI characters as collection names. You cannot use special characters or spaces.
  - **Add**: Click the Add icon to add the RPZ and collection name to the table.
  - **Save and Close**: Click this to save the configuration and close the editor.

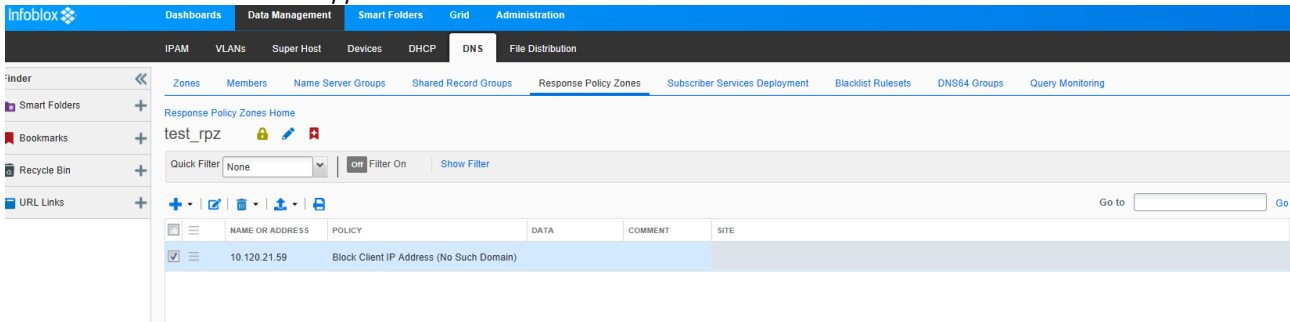
You can do the following in this tab:

- To edit an entry in the list, click the checkbox beside an RPZ, and then click the Edit icon.
- To delete an entry in the list, select the checkbox beside an RPZ, and then click the Delete icon.

## Mapping RPZs with TAXII Collection



## RPZ Rules created for the Mapped RPZ and Collection



## Starting and Stopping the TAXII Service

To start the TAXII service:

1. From the **Grid** tab, select the **Services** tab -> **TAXII\_member** checkbox and then click the Start icon from the vertical Toolbar.

To stop the TAXII service:

1. From the **Grid** tab, select the **Services** tab -> **TAXII\_member** checkbox and then click the Stop icon from the vertical Toolbar.

## Extensible Attributes for TAXII Service

You can define extensible attributes that are specific to the TAXII service. When you define TAXII specific extensible attributes, the RPZ rules created will have these attributes and their corresponding values (received in the TAXII messages) added automatically.

For information about how to configure extensible attributes, see [Managing Extensible Attributes](#).

### Extensible attributes for TAXII service

Attribute Name	Attribute Type	Description
TAXII_collection	String	The name of the TAXII collection the TAXII client delivered the message to.
TAXII_source	String	The IP address of the TAXII client that sent the TAXII message.
TAXII_member	String	The TAXII Grid member that receives TAXII message resulting in the creation of the RPZ rule.
TAXII_timestamp	Date/Integer	The timestamp when the TAXII message was received.
TAXII_user	String	The login name of the user the TAXII client connected as to the TAXII server on the member that received the message.

### Monitoring TAXII Server

You can monitor the status of the TAXII server, as described in [Monitoring Grid Services](#). If there are any invalid TAXII messages, the appliance makes a syslog entry. For information, see [Verifying RPZ in the Syslog](#). The appliance also sends an SNMP trap and an email notification, if configured. For information about setting SNMP and email notification, see [Configuring SNMP](#).

### Managing RPZs

You can manage RPZs that you defined earlier and modify their information. You can perform the following:

- [Viewing RPZs](#)
- [Modifying RPZs](#)
- [Reordering RPZs](#)
- [Locking and Unlocking RPZs](#)
- [Deleting RPZs](#)

### Viewing RPZs

You can view the list of RPZs, local, feed, or FireEye integrated RPZs, which are currently listed in the Grid. To view RPZs, complete the following:

1. From the **Data Management** tab -> **DNS** tab -> **Response Policy Zones** tab.
2. Grid Manager displays the following:
  - **Order**: Displays the order of RPZs. The order value is empty if you do not assign a primary name server when configuring a local RPZ, or if the local RPZ or the service is disabled.
  - **Name**: Displays the name of the RPZs. Click the RPZ link to view the following details:
    - **Name or Address**: Displays the domain name or the IP address.
    - **Policy**: Defines the policy defined for the corresponding domain name or IP address.
    - **Data**: Displays the target data of the rule.
    - **Comment**: Displays the comment specified when an RPZ is defined.
    - **Disabled**: Displays **Yes** if the RPZ rule is disabled.
    - **Site**: Displays extensible attributes that are associated with the domain name or IP address.
  - **Type**: Displays the type of RPZs, that is, **Local**, **Feed**, or **FireEye**.
  - **Primary Name Server**: Displays the primary name server that is associated with an RPZ.

- **Last Updated:** Displays the last updated time. For RPZ feed, it indicates if the RPZ feed has stalled and when the last zone transfer happened. For a local and FireEye integrated RPZ, it indicates the last time the zone or data was modified.
- The last updated time is empty, if:
  - A local RPZ is not associated with a primary Grid name server.
  - A zone, either a local RPZ or an RPZ feed, is not enabled.
  - An inbound zone transfer has not occurred for an RPZ feed.
  - Member's DNS service is disabled.
- **Comment:** Displays the comment recorded when creating the zone. You can double-click on a row to edit the comment. Click **Save** after modification.

For FireEye integrated RPZs, this column displays the comment recorded when creating the FireEye integrated RPZ. The rules that are created from the FireEye alerts will have alert information in this column. This differentiates between FireEye alert created rules and user created rules. You can double-click on a row to edit the comment. Click **Save** after modification. Infoblox recommends that you do not modify any internal objects. For example, the **Comment** column has alert related information, if you modify the data, then the actual alert data will be compromised.

- **Disabled:** Displays **Yes** if the RPZ is disabled. Otherwise, this field displays **No**.
- **Locked:** Displays **Yes** when a zone is locked by an admin, and displays **No** when the zone is unlocked.
- **Site:** Displays the values that were entered for this pre-defined attribute. You can double-click on a row to edit the Site. Click **Save** after modification.

You can also perform the following:

- Use **Quick Filter** and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches. Select a value from the drop-down list to filter the RPZs.
  - **None:** Select this to display all the RPZs that you have configured.
  - **All Local Response Policy Zones:** Select this to list only the local RPZs.
  - **All Feed Response Policy Zones:** Select this to list only the RPZ feeds.
  - **All Fire Eye Response Policy Zones:** Select this to list only the FireEye RPZs.
- Create a quick filter to save frequently used filter criteria. For more information, see [Using Quick Filters](#).
- You can create a bookmark for the RPZs. For more information, see [Using Bookmarks](#).
- You can modify some of the data in the table. Double-click a row of data, and either edit the data in the field or select an item from a drop-down list. Note that some fields are read-only. For more information, see [Modifying Data in Tables](#).
- To export the list of RPZs to a .csv file, click the *Export* icon. For more information, see [Importing and Exporting Data using CSV Import](#).
- Click the *Print* icon to print the list of RPZs. For more information, see [Printing from Grid Manager](#).

## Modifying RPZs

You can modify the name servers or name server groups, update policy override details and permissions, or edit extensible attributes that are associated with an RPZ.

To modify RPZs, complete the following:

1. From the **Data Management** tab -> **DNS** tab -> **Response Policy Zones** tab -> *Response Policy Zone* checkbox and then click the *Edit* icon.
2. The RPZ editor provides the following tabs from which you can modify data:
  - **General:** In this tab, you can change the information you previously entered through the wizard, as described in [Configuring Local RPZs](#). For FireEye integrated RPZs, you can update the policy type, comments, enable or disable, or lock the zone. For more information, see [Configuring FireEye RPZs](#).
  - For a FireEye integrated RPZ, the **FireEye** tab is displayed. This tab is displayed only after you install the Security Ecosystem license. You can modify or override the rule mapping for FireEye alerts or APT events. For more information, see [Configuring FireEye RPZs](#).

- You can also enter or edit information in the **Name Servers, Extensible Attributes, Settings, and Permissions** tabs. For more information, see [Modifying, Disabling, and Deleting Host and Resource Records](#).
- **Logging:** In this tab, you can enable or disable logging at the zone level for RPZ zones. You can **Override the RPZ logging** option only if the **RPZ/Security log** is enabled at the Grid or member level. The values are inherited from the Grid or member by default. To view RPZ logs at the Grid level, see [Setting DNS Logging Categories](#). Selecting the **Override** option allows you to disable RPZ logging for the particular zone. However, if the **RPZ logging** is not enabled at the Grid or member level, you cannot **Override** or **Inherit** logging at the zone level.

3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Reordering RPZs

You can change the order of RPZs, local feeds, or FireEye integrated RPZs, in each view. When you add a new local RPZ, it is added to the top of the zone list and an RPZ feed is automatically added to the bottom of the zone list. You can change the order of each through the re-ordering process.

The policy override works based on zone ordering. The zone at the top has the highest priority and it overrides the lower priority zone. To override an RPZ feed with a local RPZ, place the local feed at the top before an RPZ feed. You cannot reorder zones if they are disabled or do not have any primary name server assigned.

To reorder RPZs, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, click **Order Response Policy Zones** from the **Toolbar**.
2. The following are displayed in the *Order Response Policy Zones* wizard:
  - **Ordering:** Use the up and down arrows to move the RPZ to the desired order.
  - **Response Policy Zone:** Displays all the RPZs.
  - **Priority:** Displays the order of RPZs.
3. Click **OK** to save the changes.

## Locking and Unlocking RPZs

You can lock an RPZ so only you can make changes to it, which prevents others from making conflicting changes. When you lock an RPZ, the Grid Manager displays LOCKED beside the RPZ. When other administrators try to make changes to a locked RPZ, the system displays a warning message that the RPZ is locked and the name of the admin who locked the RPZ.

Only a superuser or the administrator who locked the RPZ can unlock it. RPZ locks do not expire and you must manually unlock a locked RPZ.

To lock or unlock RPZs, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, select the *Response Policy Zone* -> *Ruleset*.
2. You can do the following:
  - **To Lock:** Click the *Lock* icon to lock the zone.
  - **To Unlock:** Click the *Unlock* icon to unlock the zone.

## Deleting RPZs

You can delete RPZs or schedule them for deletion at a later date. The NIOS appliance moves the deleted RPZs to the Recycle Bin if enabled. When you restore the zone from the Recycle Bin, it will be restored to the bottom of the zone list.

To delete RPZs, complete the following:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab -> *Response Policy Zone* checkbox.

2. To delete an RPZ immediately, click the *Delete* icon, and then click **Yes** to confirm the delete request.
3. To schedule the deletion, click **Schedule Deletion** and in the *Schedule Change* panel, enter a date, time, and time zone. For more information, see [Scheduling Deletions](#).

The Grid Manager moves the RPZ to the Recycle Bin, from which you can restore it or permanently delete it.

## Managing RPZ Rules

You can manage local RPZ, including FireEye integrated RPZ rules that you defined earlier and modify their information. You can do the following:

- [Viewing RPZ Rules](#)
- [Modifying RPZ Rules](#)
- [Deleting RPZ Rules](#)
- [Copying RPZ Rules](#)
- [Importing RPZ Rules](#)
- [Disabling NSDNAME and NSIP rules for RPZ zones](#)

### Note

You cannot modify the rules of an RPZ feed. However, you can override the entire ruleset or each rule using local RPZs.

## Viewing RPZ Rules

You can view and edit the rules that are defined for each local RPZ, including FireEye integrated RPZs. To view RPZ rules:

1. From the **Data Management** tab -> **DNS** tab -> **Response Policy Zones** tab -> click the RPZ link.
2. You can view the following:
  - **Name or Address:** Displays the domain name or the IP address on which the rule is defined.
  - **Policy:** Displays the rule applied on the domain name or the IP address.
  - **Data:** Displays the target data of the rule.
  - **Comment:** Displays the comment specified when the rule is defined.
  - **Disabled:** Displays **Yes** if the RPZ rule is disabled. Otherwise, this field displays **No**.
  - **Site:** Displays an extensible attribute, **Site**.
  - **Expiration:** Displays the expiration time for the corresponding FireEye integrated RPZ rule. Note that NIOS updates the expiration time when you change the lifetime of the FireEye integrated RPZ rule, or if the last updated time of the rule changes, or if the alert type that generates the rule changes. This time is estimated based on the following:  
Expiration Time = Lifetime of an alert type + Last updated time of the rule  
NIOS runs a scheduler every 10 minutes to identify FireEye integrated RPZ rules whose expiration time is less than the current time. If there are rules whose expiration time is less than the current time, then such rules will be deleted. NIOS logs all deletion activities in the syslog. You can view the syslog to verify expired rules. For more information, see [Viewing RPZ in the Syslog](#).
  - **Fire Eye Alert Type:** Displays the type of FireEye alert.
  - **Last Updated:** Displays the time when the RPZ rule was last updated.

### Note

The columns, **Expiration**, **FireEyeAlertType**, and **LastUpdated**, are displayed only for FireEye integrated RPZ rules. These columns are not displayed for non-FireEye RPZ rules.

You can also do the following:



- Use **Quick Filter** and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria. For information, see [Using Quick Filters](#).
- Modify some of the data in the table. Double click a row of data, and either edit the data in the field or select an item from a drop-down list. You can edit **Comments** and **Extensible Attributes**. Note that some fields are read-only. For more information about this feature, see [Modifying Data in Tables](#).
- To export the list of RPZ rules to a .csv file, click the *Export* icon. For information on the export options, see [Importing and Exporting Data using CSV Import](#).
- Click the *Print* icon to print the list of RPZ rules. For more information, see [Printing from Grid Manager](#).

## Modifying RPZ Rules

You can modify the name of a local or FireEye integrated RPZ rule, IP address, network address, substituted name, and the comment recorded for the corresponding rule. You can also update the TTL settings or the extensible attributes that are associated with an RPZ rule.

To modify RPZ rules:

1. From the **Data Management** tab -> **DNS** tab -> **Response Policy Zones** tab -> click the RPZ link -> *Name or Address* checkbox, and then click the *Edit* icon.
2. The RPZ rules editor provides the following tabs from which you can modify data:
  - In the **General** tab, you can change the information you previously entered through the wizard. For more information, see [Configuring Rules for RPZs](#).
  - You can also enter or edit information in the **TTL and Extensible Attributes** tabs. For information about TTL settings, see [Specifying Time To Live Settings](#). For information about extensible attributes, see [Managing Extensible Attributes](#).
3. Save the configuration.

## Deleting RPZ Rules

You can delete local RPZ rules, including FireEye integrated RPZ rules, or schedule them for deletion for a later date. When you remove an RPZ rule, the NIOS appliance moves it to the Recycle Bin, if enabled.

To delete RPZ rules:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab -> *Response Policy Zone -> Ruleset*.
2. To delete an RPZ rule immediately, click the *Delete* icon, and then click **Yes** to confirm the delete request. To schedule the deletion, click **Schedule Deletion** and in the **Schedule Change** panel, enter a date, time, and time zone. For information, see [Scheduling Deletions](#).

Grid Manager moves the RPZ rule to the Recycle Bin, from which you can restore or permanently delete it.

## Copying RPZ Rules

You can copy rules from one local RPZ to another local RPZ or from one FireEye integrated RPZ to another FireEye RPZ. You can also copy rules from a local RPZ to a FireEye integrated RPZ or vice-versa. Different views of the same RPZ may have a number of rules in common. If this is the case, you can copy rules between views and zones.

To copy RPZs between DNS zones and views:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab, click **Copy Rules** from the **Toolbar**.
2. In the *Copy Rules* dialog box, Grid Manager displays the last selected zone or the zone from which you are copying rules in the **Source** field. The following fields are displayed:
  - **Source:** Grid Manager displays the last selected zone or the zone from which you are copying rules. It also displays the associated DNS view.
  - **Destination:** Click **Select Zone** to select the destination zone. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box from which you can select one. After you select the zone, Grid Manager displays the associated DNS view.
  - **Copy All Rules:** Select this option to copy all the rules.



- **Copy Specific Rules:** Select this option to copy specific rules only. Select a rule from the **Available** column and click the right arrow to move it to the **Selected** column.
  - **Copy Options:** Select one of the following:
    - **Delete all rules in the destination before copying the rules:** Select to delete all rules in the destination zone before the records are copied.
    - **Overwrite existing rules:** Select to overwrite existing rules that have the same domain name owners as the rules being copied.
3. Click **Copy & Close**.

## Importing RPZ Rules

You can import rules from an RPZ zone to a local zone. To import, you must enable zone transfer on the external server. The rules of the existing zone are overwritten when you import rules from an external server.

To import RPZ rules:

1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab -> click the RPZ link in the **Name** column -> *Rules*, click **Import Zone** from the **Toolbar**.
2. In the *Import Zone* dialog box, the following fields are displayed:
  - **Zone:** The RPZ that you have selected is displayed.
  - **DNS View:** The DNS view that you have selected is displayed.
  - **Address:** Enter the address of the external server from where you want to import rules.
3. Click **Import**.

## Disabling NSDNAME and NSIP rules for RPZ zones

NSDNAME and NSIP rules are enabled for RPZ zones by default. You can disable or enable NSDNAME and NSIP rules to control the validation of NS records and glue records received by upstream DNS servers. When you enable RPZ on internal DNS servers and if there are forward-mapping zones that are not reachable from external networks, NSDNAME and NSIP validation is not necessary. In this case, you can disable NSDNAME and NSIP rules to reduce delays in responses. When you disable these rules for RPZ zones, the appliance bypasses NSDNAME and NSIP validation for the queries and it significantly improves the performance. Note that this setting disables both NSDNAME and NSIP rules at the same time for both internal and external RPZ zones. This setting only affects the lookup process but the zone data remains unchanged. NSDNAME and NSIP records will still be available in RPZ zones during zone transfers (AXFR and IXFR). If you disable NSDNAME and NSIP rules for RPZ zones at the Grid level, all members inherit this setting. You can override this setting for each member.



### Note

Disabling NSDNAME and NSIP rules for RPZ zones on members which may send recursive queries to external servers, results in reduced security. Accordingly, it is not recommended to disable these rules for RPZ zones on members that respond with data from external servers.

To disable NSDNAME and NSIP rules for RPZ zones:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, expand the **Toolbar** and click **Grid DNS Properties**.  
**Member:** From the **Data Management** tab, select the **DNS** tab, click the **Members** tab -> *member* checkbox -> Edit icon.
2. In the *Grid DNS Properties* or *Member DNS Properties* editor, click the **General** tab -> **Advanced** tab and complete the following.
  - **Disable NSDNAME and NSIP rules for RPZ zones:** This checkbox is deselected and NSDNAME and NSIP rules are enabled for RPZ zones by default. Select this checkbox to disable NSDNAME and NSIP rules for all RPZ zones. To override the value inherited from the Grid, click **Override**. To retain the same value as the Grid, click **Inherit**.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring Prefix Length Limit for RPZ-IP Triggers

To avoid the possibility of DNS outage resulting from errors in the RPZ rules received from external sources by the RPZ feed or due to errors in the RPZ rules added to local RPZ, Infoblox provides an option to set the prefix length limit for RPZ-IP triggers. This enables the appliance to ignore RPZ-IP rules with prefix lengths that are less than the configured minimum prefix length, and to enforce only those RPZ-IP rules whose prefix lengths are equal to or greater than the configured minimum prefix length, thus accepting legitimate queries instead of dropping all queries. For example, if you configure 24 as the minimum IPv4 prefix length, the Grid enforces only those RPZ-IP rules with prefix length equal to or greater than 24 and the RPZ-IP rules with prefix lengths less than 24 are not enforced on queries that originate from external sources.

You can configure the prefix length limit for IPv4 and IPv6 prefixes at the Grid level and override it for a member, DNS view, or RPZ zone. The appliance logs a warning message in the syslog when RPZ-IP rules with prefix length less than the configured minimum prefix length are added to the local RPZ and, when an RPZ feed receives RPZ-IP rules with prefix length less than the configured prefix length from external sources.

To configure the prefix length limit for RPZ-IP triggers:

1. **Grid:** From the **Data Management** tab, select the **DNS** tab, and then select **Grid DNS Properties** from the Toolbar.  
**Member:** From the **Data Management** tab, select the **DNS** tab -> **Members** tab -> *member* checkbox -> Edit icon.  
**DNS View:** From the **Data Management** tab, select the **DNS** tab and click the **Zones** tab -> *dns\_view* checkbox -> Edit icon.  
**RPZ Zone:** From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab -> *Response policy zone* checkbox -> Edit icon.
2. In the editor, click **Toggle Advanced Mode**, and then select the **Security** tab.
3. Complete the following in the Response Policy Zones section of the **Security** tab:  
To override the Grid settings, click **Override**. To retain the same settings as the Grid, click **Inherit**.
  - **Ignore RPZ-IP triggers with too small prefix lengths:** Select this checkbox to set the prefix length limit for RPZ-IP triggers and enable the appliance to ignore the RPZ-IP rules with prefix lengths that are less than the specified prefix length limit. This checkbox is deselected by default.
    - **Minimum IPv4 Prefix Length:** Enter the minimum prefix length for IPv4 prefixes. You can specify a value between 1 to 31. The default value is 29.
    - **Minimum IPv6 Prefix Length:** Enter the minimum prefix length for IPv6 prefixes. You can specify a value between 1 to 127. The default value is 112.
4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Configuring Thresholds for RPZ Hit Rate

When the RPZ hit rate, which is the ratio of the number of queries that result in modifying the genuine response due to RPZ rules to the total number of incoming queries is high, it is unexpected and might warrant your attention. Note that the queries that hit passthru RPZ rules are not considered for the RPZ hit rate. You can configure thresholds for RPZ hit rate, above which the appliance makes a syslog entry and sends alerts as SNMP traps and email notifications. Note that you must enable notifications in order for the appliance to send SNMP traps and email notifications. For information about setting the SNMP trap and email notifications, see [Setting SNMP and Email Notifications](#).

Note that the appliance calculates the RPZ hit rate globally for all DNS views and sometimes the RPZ hit rate might be misleading. For example, if there are multiple DNS views with or without RPZ rules, there is a possibility that some DNS views might receive a substantial number of normal queries, obscuring the possible high RPZ hit rate in the other DNS views. Also, when the DNS server is configured for both authoritative and recursive queries, it is possible that the authoritative zones receive a substantial number of queries for which RPZ rules are not considered. It might make the resulting RPZ hit rate normal even if there is an excessive number of hit for recursive queries.

To configure the thresholds for RPZ hit rate:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.  
**Member:** From the **Grid** tab, select the **Grid Manager** -> **Members** tab -> *member*, and then click the **Edit** icon.
2. In the *Grid Properties* or *Grid Member Properties* editor, click **Toggle Advanced Mode**, and then select the **SNMP Threshold** tab.
3. Complete the following in the Response Policy Zones Hit Rate Configuration section of the **SNMP Threshold** tab.

- **RPZ Hit Rate:** Click **Override** to override the inherited settings, and specify the following:
  - **Trigger %:** Enter the Trigger value between 0 and 100. If the RPZ hit rate equals the Trigger value, the appliance logs a syslog entry and — if configured to do so — sends an SNMP trap and an email notification. The default Trigger value is 10%.
  - **Reset %:** Enter the Reset value between 0 and 100. If the RPZ hit rate equals the Reset value, the appliance logs a syslog entry and — if configured to do so — sends an SNMP trap and an email notification, to notify that the RPZ hit rate has gone back to an acceptable level. The default Reset value is 2%.
- **Interval:** Enter the time interval that determines when the appliance starts calculating the RPZ hit rate. You can enter a value between 1 and 86400. The default value is 10 seconds. At the end of each interval, if the number of incoming queries equals or exceeds the **Minimum query** value, the appliance calculates the RPZ hit rate and if the RPZ hit rate exceeds the **Trigger** value, the appliance sends notifications and continues to send notifications at the end of subsequent intervals, until the RPZ hit rate equals the **Reset** value.  
Note that the appliance calculates the RPZ hit rate at the end of each **Interval** or when the number of incoming queries reach the **Maximum query** value, whichever comes sooner.
- **Minimum query:** Specify the minimum number of queries received between the RPZ hit rate checks. The default value is 1000. The appliance calculates the RPZ hit rate when the number of incoming queries equals or exceeds the **Minimum query** value at the end of the **Interval**. If the total number of incoming queries is less than the **Minimum query** value, the appliance skips the RPZ hit rate check and the query count continues to cumulate into subsequent intervals until the **Minimum query** is met.
- **Maximum query:** Specify the maximum number of queries received between the RPZ hit rate checks. The default value is 100000. When the number of incoming queries equals or exceeds this value, the appliance calculates the RPZ hit rate and does not wait for the expiration of the **Interval**.

4. Save the configuration and click **Restart** if it appears at the top of the screen.

## Verifying RPZ Configuration

After you have set up and configured RPZs and RPZ rules, you can verify whether the RPZ zone transfers are functioning properly by doing the following:

- [Viewing RPZ in the Syslog](#)
- [Viewing the Last Updated RPZs](#)

The appliance also makes a syslog entry, when an RPZ zone refresh succeeds or fails and also sends an SNMP trap and an email notification, if configured. For information about setting SNMP and email notification, see [Setting SNMP and Email Notifications](#).

### Viewing RPZ in the Syslog

To receive RPZ information in the syslog, make sure that you enable the RPZ option in the Logging tab of the Grid DNS Properties editor. Once the RPZ option is enabled, the appliance logs RPZ threats in CEF (Common Event Format) in the syslog. You can click the Action icon to view the RPZ threat details in the RPZ Threat Details viewer. For information about how to configure the syslog server, see [Using a Syslog Server](#).

Following is a sample RPZ threat message:

```
2014-09-15T07:14:47-07:00 daemon info rpz:
CEF:0|Infoblox|NIOs|6.12.0-252689|RPZ-QNAME |PASSTHRU|7|app=DNS
dst=172.31.1.156
src=10.120.20.69 spt=39503 view=_default qtype=A msg="rpz QNAME PASSTHRU
rewrite
passthru.com [ANY] via passthru.com.rpz_1.com
```

Each log message contains the following information:

- The timestamp when the event happened in yyyy-mm-ddThh:mm:ss-00:00 format.
- **Infoblox|NIO** **x.x.x**: Indicates the Infoblox product, and x.x.x represents the NIOS version.
- The string following the NIOS version is a hard-coded constant. In this example, it is RPZ QNAME.
- The hard-coded constant is followed by mitigation action. In this example, it is PASSTHRU.
- The number following the mitigation action is the threat severity level. The following numbers indicate the severity levels:
  - **8 = Critical**
  - **7 = Major**
  - **6 = Warning**
  - **4 = Informational**
- **dst**: Destination IP address.
- **src**: Source IP address.
- **spt**: Source port.
- **view**: DNS view.
- **qtype**: Query type.
- **msg**: RPZ rule.

The syslog messages are optionally tagged according to the logging category configured in the external syslog servers.

For more information, see [Syslog Message Prefixes](#).

To verify RPZ zone transfers:

1. Go to the **Administration** tab -> **Logs** tab -> **Syslog** tab.
2. Select **RPZ Incident Logs** from the **Quick Filter** drop-down list.
3. Review the syslog for zone transfer confirmation, as shown in the below figure.

### The Syslog Viewer

Sample RPZ Threat Message

The screenshot shows the Infoblox Syslog Viewer interface. The top navigation bar includes 'Administration', 'Logs', and 'Syslog'. The 'Quick Filter' is set to 'RPZ Incident'. The main area displays a table of log messages with columns for 'TIMESTAMP', 'FACILITY', 'LEVEL', 'SERVER', and 'MESSAGE'. A red box highlights a message with a severity level of 8 (Critical) and a mitigation action of 'PASSTHRU'.

TIMESTAMP	FACILITY	LEVEL	SERVER	MESSAGE
2018-12-21 16:...	daemon	INFO	named[10389]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=53709 view=_default qtype=A msg='rpz...
2018-12-21 16:...	daemon	INFO	named[10389]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=47404 view=_default qtype=A msg='rpz...
2018-12-21 15:...	daemon	INFO	named[10389]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=49269 view=_default qtype=A msg='rpz...
2018-12-21 15:...	daemon	INFO	named[10389]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=36971 view=_default qtype=A msg='rpz...
2018-12-21 15:...	daemon	INFO	named[10389]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=49328 view=_default qtype=A msg='rpz...
2018-12-21 15:...	daemon	INFO	named[21909]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=57143 view=_default qtype=A msg='rpz...
2018-12-21 15:...	daemon	INFO	named[21909]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=58527 view=_default qtype=A msg='rpz...
2018-12-21 15:...	daemon	INFO	named[21909]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=42195 view=_default qtype=A msg='rpz...
2018-12-21 15:...	daemon	INFO	named[21909]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=60740 view=_default qtype=A msg='rpz...
2018-12-21 15:...	daemon	INFO	named[21909]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=51670 view=_default qtype=A msg='rpz...
2018-12-21 15:...	daemon	INFO	named[21909]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=35148 view=_default qtype=A msg='rpz...
2018-12-21 15:...	daemon	INFO	named[21909]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=44330 view=_default qtype=A msg='rpz...
2018-12-21 15:...	daemon	INFO	named[21909]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=39657 view=_default qtype=A msg='rpz...
2018-12-21 15:...	daemon	INFO	named[21909]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=48952 view=_default qtype=A msg='rpz...
2018-12-21 15:...	daemon	INFO	named[21909]	rpz: CEF:0 Infoblox NIO 8.4.0-378940 RPZ-CLIENT-IP INXDOMAIN 7 app=DNS dst=10.35.2.86 src=10.120.21.59 spt=57388 view=_default qtype=A msg='rpz...

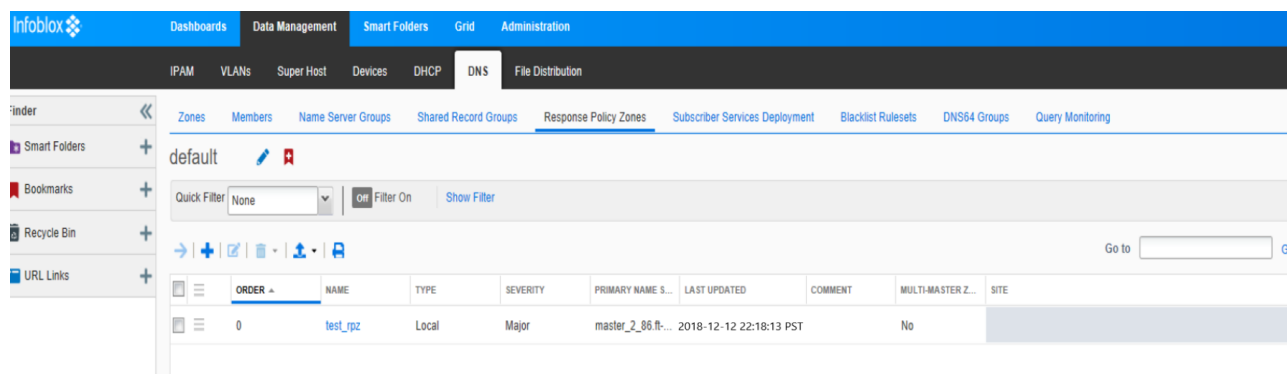
### Viewing the Last Updated RPZs

To view the last updated RPZs:

1. Go to the **Data Management** tab -> **DNS** tab -> **Response Policy Zones** tab.
2. Review the **Last Updated** column and confirm the time when an RPZ was last updated, as shown in the following figure.

 It may take up to 10 minutes before the updated information is displayed.

### Last Updated RPZ



ORDER	NAME	TYPE	SEVERITY	PRIMARY NAME S...	LAST UPDATED	COMMENT	MULTI-MASTER Z...	SITE
0	test_rpz	Local	Major	master_2_86.ft...	2018-12-12 22:18:13 PST		No	

## Best Practices for Configuring RPZs

Before configuring RPZs, observe the following best practices to ensure a successful configuration:

### General RPZ Best Practices

- When you enable Infoblox DNS Firewall, DNS performance for all queries, recursive or authoritative, will be affected.
- For performance reasons, Infoblox recommends that you maintain a reasonable number of zones.
- Do not enable RPZ on multiple layers, such as on DNS client facing servers and forwarders.
- If you have multiple DNS servers in a Grid, ensure that you configure RPZs on the recursive server that is closest to your DNS clients. If you configure RPZs on second level DNS caching servers, you will not be able to identify the DNS clients because only the IP addresses of the forwarding name servers can be identified.
- Infoblox recommends that you preview your RPZ rules to ensure ruleset integrity and to avoid unexpected results. You can preview your rules by selecting **Log Only (Disabled)** when you configure **Policy Override** for an RPZ, RPZ feed, or FireEye integrated RPZ. For information about how to configure this, see [Configuring Local RPZs](#) and [About FireEye Integrated RPZs](#).
- The appliance logs all matching and disabled rules for all queries in the syslog. You can view the syslog to ensure that the rules are set up correctly before they take effect. Ensure that you enable **rpz** in the **Logging Category** of *Grid DNS Properties* editor to log these events. For information about how to set logging categories as described in setting DNS logging categories, see [Using a Syslog Server](#).
- You can use the standard TSIG mechanism to ensure that feed zones come from the correct servers. Grid members can function either as a primary or secondary servers for the RPZ. As with hosting any zone as a secondary, please ensure that the appliance is sized properly to hold the zone contents in memory.
- You can only export or import the RPZ local zones using the CSV export or import feature, but you cannot import or export FireEye zones using this feature.
- Note that the NIOS blacklist and NXDOMAIN features take precedence over RPZs.
- In order to leverage DNS notify messages to trigger zone transfer of the feed zone, port 53 of the lead secondary must be open to receive such messages. If not, the zone will refresh based on the refresh setting in the SOA.
- The name of the zone, which is assigned to an RPZ member, must not exceed 241 characters. When the name exceeds this limit, respective zone fails to load.

- For RPZs that contain IP addresses, RPZ query name recursion continues to take place irrespective of other settings (such as the **Enable RPZ query name recursion (qname-wait-recurse)** checkbox). Recursion takes place for the first RPZ that contains the IP rules. Because of this, if data exfiltration through DNS is to be blocked, then RPZs associated with disrupting the blockage must be placed before RPZs containing the IP rules. That is, RPZs that use IP rules must be placed last in the RPZ order.

## Best Practices For FireEye Integrated RPZs

Before you configure a FireEye integrated RPZ, consider the following:

- FireEye integrated RPZs inherit default values from local RPZs. You can create, edit and delete rules using the Infoblox GUI, API, and RESTful API.
- To avoid false positives, Infoblox recommends that you create a whitelist of allowed zones using a local RPZ that is sorted above the FireEye RPZ and add your own domain to the whitelist RPZ. For example, you can add your company domain name, such as corpxyz.com. This list must contain popular domains, such as Alexa 250, and other desired domains.
- Note that there will be an impact on the storage capacity when you create a new FireEye alert and map it with an RPZ rule. The processing of alerts will consume a few CPU cycles, which will have some impact on the system.
- You must properly configure the settings on a FireEye appliance. NIOS supports only **Per Event** delivery mechanism and **JSON Normal** message format. To ensure that the NIOS appliance process alerts properly, configure the FireEye appliance accordingly. For more information about alerts as described in handling alerts from the fireeye appliance, see [About FireEye Integrated RPZs](#).
- You cannot add a FireEye integrated RPZ during a scheduled full upgrade. However, updates to the CNAME record are processed during a full upgrade. NIOS updates CNAME records in the database to store information that is specific to FireEye alerts.
- The rules created due to insertion of alerts will be visible through the FireEye RPZ viewer. Infoblox recommends that you do not modify any internal objects. For more information about viewing RPZs, see [Managing RPZs](#).
- Note that SSL certificate validation is not supported.
- You must verify the following after you configure the FireEye and NIOS appliances:
  - The URL configured on the FireEye appliance matches the URL in the FireEye integrated RPZ on NIOS.
  - Verify the username and password for FireEye admin on the FireEye appliance.
  - Ensure that the settings are properly configured on the FireEye appliance.
  - Verify the state of the FireEye appliance.

For more information about configuring the FireEye appliance to send alerts to the NIOS appliance, see [About FireEye Integrated RPZs](#).

- Note that the frequency of alerts received from FireEye can be minimal. A very small number of alerts are generated on a weekly basis. For example, the FireEye appliance may generate only tens of alerts per day.

## Configuring Combination Threat Feeds

Infoblox has a large and robust offering of threat intelligence feeds that allow you to tailor your security to your needs. In an effort to simplify the process, Infoblox has created “combination feeds” to group sets of threat intelligence feeds in order to reduce the number of choices and make the process more intuitive. As no two users have the same requirements, threat intelligence feeds have been narrowed to a small number of “sets”. The set you choose depends on the sensitivity your environment has to protection of potential threats as compared to the sensitivity your environment has to blocking potentially benign sites.

Aside from simplifying the choice from more than twenty individual feeds to a single set, there are other benefits to these feeds as well. These feeds will continue to be curated over time. As new feeds that qualify for inclusion are introduced, new classes of indicators will be added automatically and do not require a major change to your policy rules. Similarly, as some feeds are deprecated, they can be removed from the combination feed automatically, thus not requiring maintenance of your policy rules. As these feeds are enhanced and maintained, you will be updated about any substantive changes to their contents.

Infoblox provides four sets of combination feeds:

- **Low:** Blocks the fewest number of threats but also minimizes the potential of blocking benign sites. Examples of these environments may be universities, service providers, and public wifi access points.

- **Medium:** An ideal balance between detection while minimizing the potential for positives. The “Medium” set has been designed to be appropriate for most enterprise organizations. If you are unsure of which set to use, “Medium” is probably the best fit for your organization
- **High:** Designed for environments where security is the most important factor. These feeds are most appropriate for environments where communication is well understood and security of the devices is critical. Examples of environments where “High” is most appropriate include server farms, networks with IoT devices or Point-of-Sale terminals. It is not recommended for networks in which users typically surf the web or check their email.
- **Extreme:** Created to provide the greatest degree of security, but these sets are not recommended for most users as the potential for positives is much higher than normal. Use this feed at your own risk.

Each set includes two separate feeds that must be deployed together and without any other set:

- **Block:** Deploy this file with the policy action of “Block”. This should be one of the first actions in the policy list, possibly following a global allow list.
- **Log:** Deploy this file with the policy action of “Allow” (preferably with log). This should be one of the last actions in the policy list.

For example, this is what a minimal policy should look like if the “Medium” set best describes your organization.

Order	Object	Action
1	Local RPZ: Allowlist	Passthru
2	RPZ: ib-med-block	Block
3	RPZ: ib-med-log	Log Only (Disabled)

## Contents of Combination Feeds

Combination feeds are combinations of existing feeds. If you deploy either the Low, Medium, High, or Extreme combination feed set, these feeds are already combinations of one or more of the following existing feeds:

- AntiMalware
- AntiMalware\_IP
- Base Hostnames
- Bogon
- Cryptocurrency hostnames and domains
- DoH Public Hostnames
- DoH Public IPs
- Exploit Kit IPs
- Extended Base & anti-malware Hostnames
- Extended Exploit Kits IPs
- Extended malware IPs
- Extended Ransomware IPs
- Extended TOR Exit Node IPs
- Malware DGA hostnames
- Malware IPs
- Ransomware
- SURBL Fresh domains
- SURBL Multi domains
- SURBL Multi Lite domains
- TOR Exit Node IPs

This means that if you use one of the combination feed sets, you must not:

- Use it in combination with any other set. That is, do not deploy Low, Medium, and High together. Each is completely self-contained, and deploying them together will almost certainly not provide the results you are looking for.



- Use it in combination with any of the other above feeds, as these feeds are already reselected in the contents of the combination feed. Therefore, deploying them together will at minimum cause your NIOS appliance to expend more resources than necessary, and may also result in undesired actions if there are conflicting policies.

### Current Contents of Each of the Combination Threat Feeds

Feeds	Extreme Block	Extreme Log	High Block	High Log	Med Block	Med Log	Low Block	Low Log
AntiMalware	✓		✓			✓		✓
AntiMalware_IP	✓		✓			✓		
Base Hostnames	✓		✓		✓		✓	
Bogon	✓		✓		✓			✓
Cryptocurrency hostnames and domains	✓		✓			✓		✓
DoH Public Hostnames	✓		✓			✓		✓
DoH Public IPs		✓		✓		✓	✓	✓
Extended Base & anti-malware Hostnames	✓		✓		✓		✓	
Extended malware IPs	✓			✓		✓		✓
Extended Ransomware IPs	✓			✓		✓		✓
Extended TOR Exit Node IPs	✓		✓			✓		✓
Malware DGA hostnames	✓		✓		✓		✓	
Malware IPs	✓			✓		✓		✓
Ransomware	✓		✓		✓		✓	

### Best Practices for Deploying Combination Threat Feeds

Before deploying combination threat feeds, observe the following best practices to ensure a successful deployment:

1. Choose the set of feeds that best describes your organization: Low, Medium or High. If in doubt, most organizations are best served with Medium.
2. Deploy the set of feeds together. For example, if “Medium” best describes your organization, deploy both *ib-med-block* and *ib-med-log* together.
3. Do not deploy one set of threat intelligence with any other set. That is, do not deploy Medium together with High. Choose one set and deploy it together and without any other set.



4. Do not deploy a set with any feeds listed in [Contents of Combination Feeds](#). These feeds are already reflected in the combination feeds, so using them together may cause duplication, redundancy, or potential of conflict.
5. Deploy blocking feeds as close to the first policy action as possible (potentially after a global allow list).
6. Deploy logging feeds after all blocking feeds so as to not accidentally allow an indicator that another feed wanted to block.

What to do if you have already deployed a combination threat feed?

If you have already deployed a combination threat feed, review the best practice recommendations above and make sure that you have followed the guidelines described above. Specifically:

- If you have deployed two or more sets of combination feeds, **remove the lower sets** as these are already reflected in the higher sets. For example, if you have deployed all three (Low, Medium, and High), remove Low and Medium because the indicators in these sets are already reflected in High.
- If you have deployed one of the sets along with any of the feeds reflected in the Contents of Combination Feeds table at [Configuring Combination Threat Feeds](#), **remove those feeds** from the policy actions list as they are already reflected in your selected policy group.

## Infoblox Threat Insight

This section provides information about the Infoblox Threat Insight (also referred to as Threat Analytics in Grid Manager), which protects mission-critical DNS infrastructure from data exfiltration through DNS tunneling. This section includes the following topics:

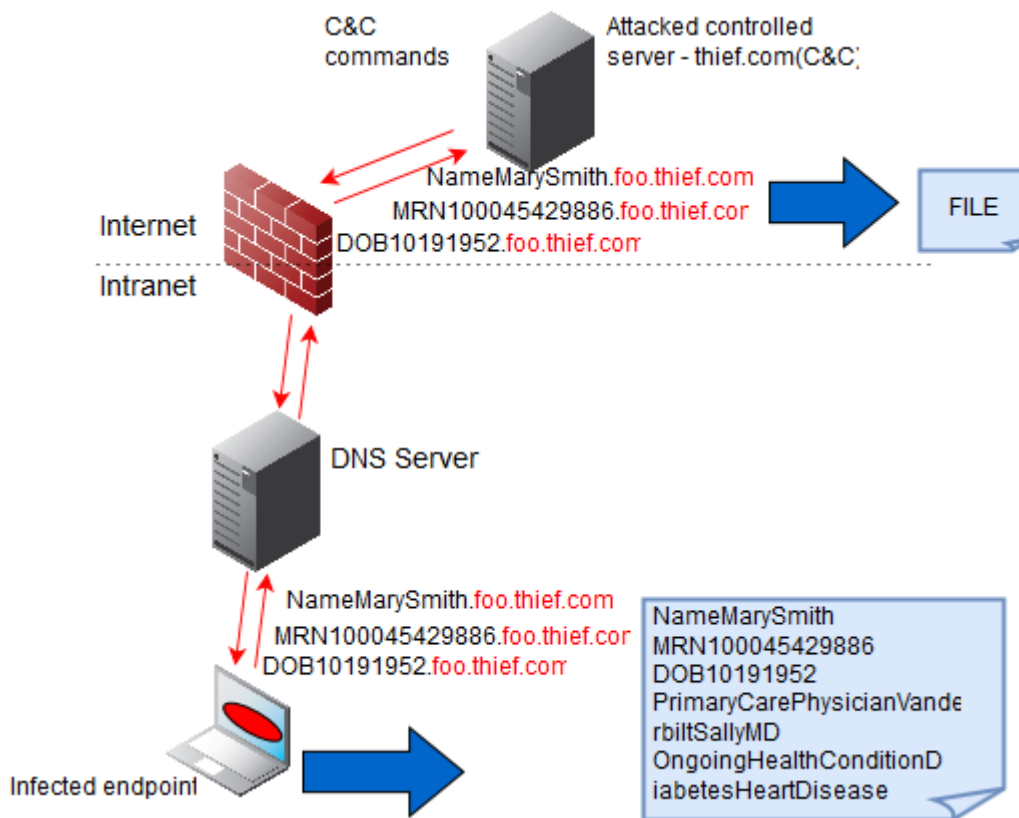
- [About Data Exfiltration](#)
- [About Infoblox Threat Insight](#)

### About Data Exfiltration

The DNS protocol is increasingly used as a pathway for data exfiltration through DNS tunneling attacks. DNS tunneling involves tunneling another protocol through port 53 — often not inspected by firewalls (even the next-generation firewalls) — by malware-infected devices or malicious insiders. There are a number of tools available for tunneling over DNS for a common motivation of bypassing captive portals for paid Wi-Fi access. A free tunneling application released under the ISC license for forwarding IPv4 traffic through DNS servers is one example of the software used in this kind of attack.

As illustrated in the below figure, sensitive information such as credit card numbers and company financial can be stolen either by establishing a DNS tunnel from within the network or by encrypting and embedding chunks of that data in DNS queries. Data is decrypted at the other end and put back together so valuable information can be stolen and misused by malicious attackers.

*Data Exfiltration*



You can use the following features to specifically target DNS tunneling traffic and minimize the risk of DNS data exfiltration:

- Anti-DNS tunneling threat protection rules: These rules detect signature-based payload encoding techniques, such as Base32, Base64 and suspicious label lengths, commonly used by tunneling products such as OzymanDNS, Iodine, DNS2TCP, and SplitBrain. For more information about the threat protection rules, refer to the Threat Protection Rules document available on the Infoblox Support site.
- Infoblox Threat Insight: Infoblox employs streaming analytics to study DNS statistics and create algorithms to identify DNS tunneling traffic. To further defend your system against DNS data exfiltration, Infoblox Threat Insight detects and mitigates DNS tunneling traffic by analyzing DNS queries and responses. Infoblox Threat Insight constantly evaluates incoming DNS traffic and develops a systematic pattern in defending against illegitimate DNS tunneling traffic and constantly updating blacklists of known malicious destinations. For information about how to configure Infoblox Threat Insight, see [About Infoblox Threat Insight](#).

## About Infoblox Threat Insight

To mitigate DNS data exfiltration, Infoblox Threat Insight (also referred to as Threat Analytics in the Infoblox GUI or Grid Manager) employs analytics algorithms to detect DNS tunneling traffic by analyzing incoming DNS queries and responses. These algorithms are developed through an extensive study and analysis of sample DNS statistics within which DNS tunneling data is identified by algorithms that cannot be detected by normal rules and signatures. For more information about DNS data exfiltration, see [About Data Exfiltration](#).

Infoblox Threat Insight identifies data exfiltration tunnels that bypass typical firewall systems. Some popular tunneling tools are OzymanDNS, SplitBrain, Iodine, DNS2TCP, TCP-Over-DNS, and others. These types of DNS threats are identified as having high activities by using the TXT records in DNS queries. Infoblox Threat Insight also identifies tunnels that are used for C&C. These threats typically do not exhibit high activities or payloads. In general, NXDOMAIN responses fall into this category of threats.

You must have at least one **RPZ** license installed in your Grid (it can be installed on any Grid member) and the **Threat**

**Analytics** license installed on the Grid member on which you want to start the threat analytics service. To download updates for threat analytics module and whitelist sets, you must have at least one **Threat Analytics** license installed in the Grid. When you enable the threat analytics service, NIOS starts analyzing incoming DNS data and applying these algorithms to detect security threats that have the same or similar behavior as the known data. Once security threats are detected, NIOS blacklists the domains and transfers them to the designated mitigation RPZ (Response Policy Zone), and traffic from the offending domains is blocked and no DNS lookups are allowed for these domains from NIOS members on which RPZ are assigned to them. The appliance also sends an SNMP trap each time it detects a new blacklisted domain.

Infoblox Threat Insight also includes a whitelist that contains trusted domains on which NIOS allows DNS traffic. These are known good domains that carry legitimate DNS tunneling traffic such as Avast, Sophos, McAfee, Boingo, Barracuda, and others. The whitelist is extensible so new whitelisted domains can be added and rolled out accordingly. For Threat Insight running on an On-Prem Infoblox DDI appliances, internal governance and vetting applied by Infoblox ensures all whitelist entries are accurate and curated, and contain only valid entries.

You can also add custom whitelisted domains or move blacklisted domains to the whitelist. For more information about how to configure Infoblox Threat Insight, see [Configuring Infoblox Threat Insight](#) below. Before you utilize Infoblox Threat Insight, there are a few guidelines you might need to consider. For more information about [Guidelines for Using Infoblox Threat Insight](#), see below.

Infoblox Threat Insight came installed with a module set and a whitelist set. To receive subsequent module set and whitelist set updates, you can configure the appliance to automatically download and apply the updates for you, or you can manually upload the updates when the appliance displays a banner message notifying about available updates. For information about how to configure the update policy, see [Defining the Threat Analytics Update Policy](#) below.

## Licensing Requirements and Admin Permissions

You must obtain and install valid licenses on your appliance before using Infoblox Threat Insight. Contact your Infoblox representative to obtain these licenses. For more information, see [Managing Licenses](#).

### Infoblox Threat Insight

To start the threat analytics service, you must have at least one **RPZ** license installed in your Grid (it can be installed on any Grid member) and the **Threat Analytics** license installed on the Grid member on which you want to start the threat analytics service. To download updates for threat analytics module and whitelist sets, you must have at least one **Threat Analytics** license installed in the Grid.

Note that running the threat analytics service might affect your system performance if the appliance has a small capacity and is taking on heavy traffic. Evaluate your Grid and Grid members to ensure that you select an appliance that is appropriate for running the threat analytics service. For more information about supported appliances, see [Supported Appliances for Infoblox Threat Insight](#) below.

### Admin Permissions

Superusers can configure all threat protection and analytics related tasks. You can assign **Security Permissions** to specific admin groups and roles so these users can configure security related tasks. You can also add a global permission for managing Grid security properties or add an object permission for managing member security properties.

To manage the analytics related tasks, you must assign appropriate read-only or read/write **Analytics Permissions** to the specified admin groups and roles. You can also add the **Global Analytics Permission** as a global permission or add **Member Analytics Permission** to specific Grid members as an object permission. For more information about how to assign admin permissions, see [Managing Permissions](#).

## Guidelines for Using Infoblox Threat Insight

Following are some guidelines to take into consideration when using Infoblox Threat Insight:

- To start the threat analytics service, you must have at least one **RPZ** license installed in your Grid (it can be installed on any Grid member) and the **Threat Analytics** license installed on the Grid member on which you want to start the threat analytics service. To download updates for threat analytics module and whitelist sets, you must have at least one Threat Analytics license installed in the Grid.

- Infoblox recommends that you run the threat analytics service for a limited time to monitor and preview what has been detected before actually blocking blacklisted domains. You can carefully review the list of detected domains and decide which domains you want to continue blocking and which domains you want to add to the analytics whitelist. You should review the blacklisted domains on a regular basis to make sure that no legitimate use of DNS tunneling is blocked. Note that you can update the analytics whitelist by adding new whitelisted domains, moving legitimate domains from the blacklisted domain list, or using CVS import and export. For more information about Configuring a Local RPZ as the Mitigation Blacklist Feed, see below.
- Analytics whitelisted domains and supported DNS tunneling tools are updated periodically and are bundled with future NIOS releases. To ensure that your appliance is using the most up-to-date whitelist, upgrade to the next NIOS release or configure the appliance to download threat analytics updates. For information about upgrades, see [Upgrading NIOS Software](#). Note that this process may change in future NIOS releases.
- There are no configurable parameters for Infoblox Threat Insight. Infoblox uses the build-in algorithms to analyze DNS statistics and blocks offending domains based on the analyzed data.
- DNS tunneling detection is not instantaneous. It may take a few seconds to a few minutes for the analytics to determine positive DNS tunneling activities.
- During an HA failover, analytics data that is in progress on the active node might be lost. Only new DNS queries on the new active node after a successful failover are being analyzed. It may take a few minutes for the analytics to reach its normal state. If there is no connection between the Grid Master and Grid member, blacklisted domains detected by the analytics cannot be transferred to the Grid Master as RPZ records for a pre-configured RPZ zone — this is not applicable to standalone appliances with RPZ license installed. In addition, ensure that the passive node must also have the RPZ license installed and that its hardware model is capable of running the threat analytics service. For information about supported appliance models, see Supported Appliances for Infoblox Threat Insight below.
- The threat analytics service only works on recursive DNS servers and forwarding servers that use BIND as the DNS resolver. It does not support Unbound as the DNS resolver.
- The analytics whitelist only applies to Infoblox Threat Insight; it does not apply to signature-based tunneling detection. Anti-DNS tunneling threat protection rules are implemented to address signature-based tunneling analysis. For detailed information about threat protection rules, refer to the *Infoblox Threat Protection Rules* available on the Support web site.
- Infoblox Threat Insight does not support RESTful APIs.

## Supported Appliances for Infoblox Threat Insight

Due to memory and capacity required to perform analytics, ensure that you install the Threat Analytics and RPZ licenses, and enable the threat analytics service on an appliance that has a big enough capacity. Following are the supported Infoblox appliance models on which you can run the threat analytics service:

- PT-1405, and PT-2205.
- IB-4015, and IB-4030-10GE.
- TE-1415, TE-1425, TE-2215, and TE-2225.
- TE-V1415, TE-V1425, TE-V2215, TE-V2225, TE-V4010, and TE-V4015.



### Note

Using unsupported appliance models for Infoblox Threat Insight might cause performance issues.

## Configuring Infoblox Threat Insight

You must have at least one RPZ license installed in your Grid (it can be installed on any Grid member) and the Threat Analytics license installed on the Grid member on which you want to start the threat analytics service. You must also create a new RPZ and use it as the designated mitigation blacklist feed so the appliance can transfer all blacklisted domains to this feed.

NIOS continuously collects and analyzes statistics of incoming queries and responses, detects possible DNS tunneling activities, blocks offending domains that match the known data, and updates the mitigation blacklist feed (a designated local RPZ) of any known malicious domains. For supported appliance models for Infoblox Threat Insight, see Supported Appliances for Infoblox Threat Insight.

To configure Infoblox Threat Insight, complete the following:

1. Obtain and install valid **RPZ** and **Threat Analytics** licenses on the appliance that is used to support analytics. Note that you must have the threat analytics service running on the member serving recursive DNS queries or have recursive DNS queries forwarded to another DNS server. To generate reports that contain statistics about DNS tunneling, you must also configure a reporting appliance in the Grid.
2. Create and add a new local RPZ and use it as the designated mitigation blacklist feed so the appliance can transfer all blacklisted domains to this feed. Ensure that you configure an appropriate policy for this RPZ. To monitor the threat analytics service before actually blocking domains, set **Policy Override** to **Log Only (Disabled)**. When you are ready to block offending domains, set **Policy Override** to **None (Given)**.
3. Configure admin permissions so admin users can manage the threat analytics service and analytics related tasks. For information about how to configure admin permission, see *About Administrative Permissions*.
4. Start the threat analytics service on the appliance that has the **Threat Analytics** license installed, as described in *Starting and Stopping the Threat Analytics Service*.



#### Note

The analytics functionality only works on recursive servers and forwarding servers that use BIND as the DNS resolver; it does not function on authoritative servers or servers that use Unbound as the DNS resolver.

After you set up Infoblox Threat Insight to mitigate DNS data exfiltration, you can do the following to manage it:

- View supported whitelisted domains for analytics, as described in *Viewing the Analytics Whitelist* below. Note that these domains are specific to analytics only. They are not used in the anti-DNS tunneling threat protection rules.
- Manually add a custom domain to the analytics whitelist, as described in *Adding Custom Whitelisted Domains* below.
- Review the blacklisted domains and make decisions about whether to move them to the analytics whitelist so future DNS activities will not be blocked. For more information, see *Viewing Blacklisted Domains* below.
- Move a blacklisted domain to the analytics whitelist, as described in *Moving Blacklisted Domains to the Whitelist*.
- Monitor DNS tunneling activities and events using pre-defined reports and the syslog, as described in *Monitoring DNS Tunneling Activities* below.

## Starting and Stopping the Threat Analytics Service

To start the threat analytics service, you must have at least one RPZ license installed in your Grid (it can be installed on any Grid member) and the Threat Analytics license installed on the Grid member on which you want to start the threat analytics service. You can also stop the service when necessary.

To start or stop the threat analytics service:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Services** tab, click the **Threat Analytics** service link. Grid Manager displays only the member or members with the RPZ license installed. Select the member checkbox.
2. From the Toolbar, click **Start** to start the service or **Stop** to stop the service.

When you stop the threat analytics service, the appliance does not detect or protect against non-signature-based DNS tunneling. In addition, reports that you generate might not include statistics related to DNS tunneling.



#### Note


After you enable the threat analytics service, you must restart DNS service for the analytics to start working.

## Viewing the Analytics Whitelist

The **DataManagement** tab -> **ThreatAnalytics** tab -> **Whitelist** tab of Grid Manager lists the trusted domains on which NIOS allows DNS traffic by default. These are known good domains that carry legitimate DNS tunneling traffic such as Avast, Sophos, McAfee, Boingo, Barracuda, and others. They are marked as **System** domains, and you cannot delete them; but you can disable them so NIOS does not treat them as trusted domains. You can also add custom domains or

move blacklisted domains to the analytics whitelist. For more information, see Adding Custom Whitelisted Domains and Moving Blacklisted Domains to the Whitelist below.

To view a complete list of trusted domains in the analytics whitelist:

1. From the **Data Management** tab, select the **Threat Analytics** tab -> **Whitelist** tab.
2. The appliance displays the following for each trusted domain:
  - **Actions:** Click the Action icon  next to a domain and select one of the following:
    - **Disable:** Click this to disable the domain. When you disable a domain, the appliance does not treat this domain as trusted domain until you enable it.
    - **Edit:** Click this to open the *Whitelist* editor. For system domains, the only property you can modify is to disable or enable them. For custom domains however, you can also add information to the **Comment** field.
    - **Delete:** This is only applicable to custom domains. You cannot delete system domains. Select this to delete the custom domain.
  - **Domain Name:** The name of the trusted domain.
  - **Type:** Displays the domain type. This can be **System** or **Custom**. A system domain is a trusted domain that carries legitimate DNS tunneling traffic such as Avast, Sophos, McAfee, Boingo, Barracuda, and others. A custom domain is one that you have added to the whitelist or moved from the mitigation blacklist RPZ.
  - **Disabled:** Indicates whether this domain is disabled or not. The appliance does not treat disabled domains as trusted domains. You can disable both system and custom domains.
  - **Comment:** Additional information about the domain.



#### Note

When you upgrade to a future NIOS release or update the analytics whitelist, all changes made to the whitelist will be preserved.

You can also do the following in this panel:

- Click **Go to Mitigation Response Policy Zone** to access the blacklisted domains that are identified as offenders for DNS tunneling. Blacklisted domains are detected through Infoblox Threat Insight and automatically transferred to the blacklist RPZ feed. For information about these domains, see Viewing Blacklisted Domains below.
- Export or import whitelisted domain names using the CSV import and export functionality.
- Navigate to the next or last page of the whitelist using the paging buttons at the bottom of the panel.
- Refresh the analytics whitelist by clicking the Refresh button.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Select a quick filter to search for **System** or **Custom** whitelist entries, or both.
- Print the whitelist or export it in CSV format.

## Adding Custom Whitelisted Domains

The analytics whitelist is populated with trusted domains that carry legitimate DNS tunneling traffic such as Avast, Sophos, McAfee, Boingo, Barracuda, and others. You can add domains that you deem trustworthy to this list. When you add a custom domain, it is marked as **Custom** in the whitelist.

To add a custom whitelisted domain, complete the following:

1. From the **Data Management** tab, select the **Threat Analytics** tab -> **Whitelist** tab, click the Add icon or click **Add Custom Whitelist** from the Toolbar.
2. In the *Add Custom Whitelist* wizard, complete the following:
  - **Domain Name:** Enter the name of the domain that you want to add to the analytics whitelist.
  - **Comment:** Enter additional information about this domain.
  - **Disable:** When you select this, the appliance does not treat this domain as a trusted domain. When you enable the domain again, it is considered as a whitelisted domain.
3. Save the configuration. You do not need to restart DNS service to update the analytics whitelist.

## Configuring a Local RPZ as the Mitigation Blacklist Feed

For the threat analytics service to function properly and for NIOS to properly report detected backlisted domains, you must create and designate local RPZs as the mitigation for the Grid. You can add any Response Policy Zones to the list of RPZs from different Network and DNS Views. When a domain is detected as malicious, NIOS will update all RPZs in the list. If you assign an existing RPZ that is used for other purposes as the mitigation blacklist feed, you may experience the following:

- Existing RPZ hits are reported as hits detected by the analytics after an upgrade.
- If you manually add rules to the RPZ, all RPZ hits are reported as hits detected by the analytics, regardless of whether they match the manually created rules or are detected through the threat analytics service.

Infoblox recommends that you run the threat analytics service for a limited time to monitor and preview what has been detected before actually blocking domains. To do so, set **Policy Override** to **Log Only (Disabled)** when you create the RPZ so you can monitor blacklisted domains without actually blocking them.



### Note

You can designate only one local RPZ as the Grid-wide mitigation blacklist feed.

To create and designate a local RPZ as the blacklist feed:

1. Create a local RPZ by completing the procedure described in [Configuring Local RPZs](#).

Note to monitor the threat analytics service without blocking domains, set **Policy Override** to **Log Only (Disabled)**. When you are ready to block blacklisted domains, set **Policy Override** to **None (Given)**.

2. From the **Data Management** tab, select the **Threat Analytics** tab -> **Whitelist** tab, click the **Grid Threat Analytics Properties** from the Toolbar.
3. In the *Grid Threat Analytics Properties* editor, click the DNS Threat Analytics tab, and complete the following:
  - Click the Add icon to open the *Zone Selector* dialog box and select the RPZs. You must configure at least one local RPZ. To remove an RPZ, select it from the table and click **Delete**.
  - Save the configuration.



### Note

You cannot delete an RPZ that is used as the mitigation blacklist feed until you remove or clear it from the *Grid Threat Analytics Properties* editor.

## Enabling Integration with BloxOne Threat Defense Cloud for Threat Insight

If your network configuration includes BloxOne Threat Defense Business On-premises, BloxOne Threat Defense Business Cloud, or BloxOne Threat Defense Advanced, you can configure a cloud integration client to collect malicious domains detected by Threat Insight in the BloxOne Threat Defense cloud. NIOS then applies the detected domains to RPZs that were configured for the on-premises Grid. This feature ensures that all malicious domains detected in BloxOne Threat Defense Cloud are also applied on Grid members on-prem.

You can use this feature when you have BloxOne Threat Defense Business On-premises, BloxOne Threat Defense Business Cloud, or BloxOne Threat Defense Advanced license. Note that you can configure only one cloud client per on-premises Grid. Ensure that you configure the email address and password in the *Grid Properties Editor* before you enable the integration with BloxOne Threat Defense Cloud Client. For more information about Configuring Integration with BloxOne Threat Defense Cloud, see [Configuring Integration with BloxOne Threat Defense Cloud](#).

To enable the integration with BloxOne Threat Defense Cloud, complete the following steps:



1. From the **Data Management** tab, select the **DNS** tab -> **Response Policy Zones** tab. Expand the Toolbar and click **BloxOne Threat Defense Cloud Client**.
2. In the *BloxOne Threat Defense Cloud Integration Client* editor, complete the following:
  - **Enable Cloud Client:** Select this checkbox to enable NIOS to get Threat Insight results in BloxOne Threat Defense Cloud.  
The results are periodically synchronized based on the interval you set. NIOS requests only subsequent data since the last data timestamp.
  - **Interval:** You can specify how often to request Threat Insight results detected in BloxOne Threat Defense Cloud in seconds or minutes. The default is 10 minutes.
  - **The list of Response Policy Zones to use for blacklisted domains:** Click the Add icon to add an RPZ to the list. When there are multiple zones, Grid Manager displays the *Zone Selector* dialog box from which you can select one. You can add RPZs from different network and DNS views.
3. Click **Save & Close**.



#### Note

Whenever a new RPZ is added and NIOS requests Threat Insight results, Grid Manager displays a *Warning* dialog box to confirm that you wish to request all detected domains by Threat Insight in BloxOne Threat Defense Cloud. If you click **No** in the *Warning* dialog box, you can use the set

`cloud_services_portal_force_refresh` CLI command in maintenance mode and set the flag to request all domains detected in BloxOne Threat Defense Cloud.

## Viewing Blacklisted Domains

To review the list of blacklisted domains, complete the following:

1. From the **Data Management** tab, select the **Data Management** tab -> **DNS** tab -> **Response Policy Zones** tab, click the mitigation blacklist RPZ name.
2. Grid Manager displays the following for each blacklisted domain:
  - **Name or Address:** Displays the name or IP address of the blacklisted domain.
  - **Policy:** Displays the policy used to handle the responses when NIOS detected the blacklisted domain.
  - **Data:** Displays the target data about this domain.
  - **Comment:** Displays additional information about this domain.
  - **Site:** This is a pre-defined extensible attribute (if configured) that is used to indicate the location of the domain.
  - **Disable:** Indicates whether this domain is disabled or not. When the domain is disabled, the appliance does not block activities on this domain, and configuration for this domain does not change. When the domain is enabled, it is considered as a blacklisted domain and all DNS activities are blocked.

You can also do the following in the blacklisted domain panel:

- Click **Go to Analytics Whitelist View** to view the analytics whitelist. In the Whitelist panel, you can see all the trusted domains for Infoblox Threat Insight, and DNS activities are allowed on these domains.
- If you want to move a blacklisted domain to the analytics whitelist so it becomes a trusted domain, select the domain checkbox and click the Action icon next to the domain, and then select **Move to Whitelist**.
- Navigate to the next or last page of the whitelist using the paging buttons at the bottom of the panel.
- Refresh the blacklist feed by clicking the Refresh button.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Select a quick filter to search for specific entries.
- Print the blacklist or export it in CSV format.

## Moving Blacklisted Domains to the Whitelist

When the appliance detects an offending domain for possible DNS tunneling, it responds according to the policy defined in the mitigation blacklist RPZ and adds the domain to the blacklist RPZ feed. You can view all blacklisted domains and turn those you deem trustworthy into trusted domains by moving them to the analytics whitelist. Note that once you move a blacklisted domain to the whitelist, you cannot reverse the action.



To move a blacklisted domain to the analytics whitelist:

1. From the **Data Management** tab, select the **Data Management** tab -> **DNS** tab -> **Response Policy Zones** tab.
2. Select a blacklisted domain and click the Action icon next to a domain and select **Move to Whitelist**.

The appliance removes the selected domain from the blacklist and adds it to the analytics whitelist. You can click **Go to Analytics Whitelist View** to verify that the domain has been successfully moved.

## Updating Threat Analytics Module and Whitelist Sets

Infoblox periodically releases threat analytics module and whitelist sets. To ensure that you can import threat analytics updates, you must have at least one Threat Analytics license installed in the Grid. The threat analytics module set consists of the analytics application .jar file, which delivers changes and updates for DNS tunneling detection; and the whitelist set consists of updated trusted domains that carry legitimate DNS tunneling traffic. You can download updates for the module set and whitelist set independently depending on how often Infoblox rolls them out. The appliance displays the version numbers of the module set and whitelist set that your Grid is currently using. To view this information before downloading updates, see [Viewing Module and Whitelist Versions](#) below.

You can configure the appliance to automatically receive and apply the latest module set and/or whitelist set. When you define an automatic update policy, the appliance checks both the analytics module and whitelist files and automatically downloads the files that have changed. You can also configure a manual update policy in which the appliance notifies you through the message banner when there are updates available. You can then decide whether you want to apply the updates to your Grid or not. For information about how to define the update policy, see [Defining the Threat Analytics Update Policy](#) below. For information about how to perform manual updates, see [Manually Uploading Threat Analytics Updates](#) below.



### Note

Only the Grid Master receives module set and whitelist set updates. Grid member receives these updates through standard Grid replication from the Grid Master. Module and whitelist data is only replicated to Grid members that have the threat analytics service enabled (an RPZ license is required to start this service on the members). The appliance uses the port 443 (HTTPS) for downloading the module set and whitelist data updates.

Infoblox recommends that you configure the appliance to automatically receive module and whitelist updates, so your appliance receives the latest information periodically. If you prefer to manually upload updates to your Grid, ensure that you apply them frequently to receive the most updated information.

## Viewing Module and Whitelist Versions

1. On the **Data Management** tab -> **Threat Analytics** tab -> **Whitelist** tab, expand the Toolbar, and then click **Grid Threat Analytics Properties**.
2. In the *Grid Threat Analytics Properties* editor, click the **Updates** tab. This tab displays the following information:
  - **Active Whitelist Version:** Displays the version number of the threat analytics whitelist set that is currently running on the Grid.
  - **Active Module Set Version:** Displays the version number of the threat analytics module set that is currently active on the Grid.

## Defining the Threat Analytics Update Policy

You can configure the settings to obtain policy updates independently for Whitelist or module set, or for both. To configure how you want to obtain the latest threat analytics updates, complete the following:

1. On the **Data Management** tab -> **Threat Analytics** tab -> **Whitelist** tab, expand the Toolbar, and then click **Grid Threat Analytics Properties**.
2. In the *Grid Threat Analytics Properties* editor, click the **Updates** tab,
3. In the *Whitelist Updates* section, complete the following:
  - **Latest Available Whitelist:** Displays the latest whitelist that is available for download.

- **Last Checked For Updates:** Displays the timestamp when the Grid last checked for updates.
  - **Whitelist Update Policy:** When you select **Automatic**, the appliance automatically downloads the latest whitelist updates based on the default or custom schedule. The appliance checks whitelist files and automatically downloads only the files that have changed. When you select an automatic policy, latest updates are activated automatically. If you select **Manual** as the update policy, the appliance displays a banner message in Grid Manager to notify you when new updates are available. You must then decide whether to apply the updates to the Grid or not. For information about how to manually apply the updates, see *Manually Uploading Threat Analytics Updates* below.
  - **Enable Automatic Whitelist Updates:** Select this checkbox to enable the automatic upload feature. When necessary, you can click **Download Whitelist Now** to override the automatic update policy. In the *Schedule* section, set up a recurring schedule for automatic updates as described in step 5.
4. In the *Module Set Updates* section, complete the following:
    - **Latest Available Module Set:** Displays the latest module set that is available for download.
    - **Last Checked For Updates:** Displays the timestamp when the Grid last checked for updates.
    - **Module Set Update Policy:** When you select **Automatic**, the appliance automatically downloads the latest module set and/or whitelist set based on the default or custom schedule. The appliance checks both the module and whitelist files and automatically downloads only the files that have changed. When you select an automatic policy, the threat analytics service on the Grid members is restarted automatically to activate the latest updates. If you select **Manual** as the update policy, the appliance displays a banner message in Grid Manager to notify you when new updates are available. You must then decide whether to apply the updates to the Grid or not. For information about how to manually apply the updates, see *Manually Uploading Threat Analytics Updates* below.
    - **Enable Automatic Module Set Updates:** Select this checkbox to enable the automatic upload feature. When necessary, you can click **Download Module Set Now** to override the automatic update policy. set up a recurring schedule for automatic updates as described in step 5.
  5. In the *Schedule* section, select one of the following to set up a recurring schedule for automatic downloads:
    - **Default:** When you select this, the appliance downloads the updates between 12:00 a.m. and 6:00 a.m. local time based on the time zone configured on your appliance. The appliance automatically selects a time between this time window the first time it performs an automatic update. All subsequent updates then follow the same schedule based on the selected time.
    - **Custom:** Select this and click the calendar icon to configure a custom schedule. Based on the policy you are configuring, in the *Automatic Whitelist Updates Scheduler* or *Automatic Module Set Updates Scheduler*, you can select **Hourly**, **Daily**, **Weekly**, or **Monthly** based on how often you want to update the module set and whitelist set.

Note that the scheduled time does not indicate the exact time for the download. Downloads occur during the mid-point during of a 30-minute time frame. Therefore, the actual download can happen 15 minutes before or after the scheduled time.

- When you select **Hourly**, complete the following:
  - **Schedule every hour(s) at:** Enter the number of hours between each update instance. You can enter a value from 1 to 24.
  - **Minutes past the hour:** Enter the number of minutes past the hour. For example, enter 5 if you want to schedule the rule update five minutes after the hour.
  - **Time Zone:** Select the time zone for the scheduled time from the drop-down list.
- When you select **Daily**, you can select either **Every day** or **Every Weekday**, and then complete the following:
  - **Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
  - **Time Zone:** Select the time zone for the scheduled time from the drop-down list.
- When you select **Weekly**, complete the following:
  - **Schedule every week on:** Select any day of the week.
  - **Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
  - **Time Zone:** Select the time zone for the scheduled time from the drop-down list.
- When you select **Monthly**, complete the following:
  - **Schedule the day of the month:** Enter the day of the month and the monthly interval. For example, to schedule the rule update on the first day after every 2 months, you can enter Day 1 every 2 month(s).

- **Time:** Enter a time in hh:mm:ss AM/PM (hours:minutes:seconds AM or PM) format. You can also select a time from the drop-down list by clicking the time icon.
  - **Time Zone:** Select the time zone for the scheduled time from the drop-down list.
6. Save the configuration.

## Manually Uploading Threat Analytics Updates

When you configure a manual update policy, the appliance notifies you about newly available module set and/or whitelist set updates. You can manually upload the updated files and apply them to the Grid.

To manually upload threat analytics updates:

1. From the **Data Management** tab, select the **Threat Analytics** tab -> **Whitelist** tab, click **Updates** -> **Manual Update** from the Toolbar.
2. The *Threat Analytics Upload* dialog displays the following:
  - **Current Whitelist Version:** Displays the version of the whitelist set that is currently running on the Grid.
  - **Last Applied On:** Displays the timestamp and time zone when the last whitelist set was applied to the Grid. This field changes each time when a whitelist set is applied.
  - **Latest Available Module Set:** Displays the version string of the latest available module set. This field changes each time when the module set is updated.
  - **Last Applied On:** Displays the timestamp and time zone when the last module set was applied to the Grid. This field changes each time when a module is applied.

To upload the module set or whitelist set:

- **File:** Click **Select** to navigate to the file location, and then upload the file. The appliance displays the file name in this field. You can upload either a module set or a whitelist set. Check the current version numbers of the whitelist and module sets to verify if they have changed before uploading new files.



### Note

You can only update to a newer whitelist set even though you can switch back to an older version of module set, if any. However, if you have configured an **Automatic** update policy, the appliance overwrites the older file version with the new one. To avoid this, you can change the update policy to **Manual** or disable automatic downloads.

Click **Test** to check the changes that will occur during the update, without actually applying the update. You can view update details in the Syslog Viewer. The appliance preserves the uploaded file if you do not click **Update** to update the module set or whitelist set. When you manually upload next time, this file name is displayed in the dialog. You can then choose to apply the update from this file or upload a new file before performing the update. Uploading a new file will remove the file that has not been applied.

3. Click **Update** to update the module set or whitelist set. You can also click **View Update Results** to view the update results.

## Monitoring DNS Tunneling Activities

You can monitor DNS tunneling activities through the following:

- **Pre-defined Reports:** If you have a reporting appliance configured in the Grid, you can generate the following reports that include DNS tunneling data.  
For more information about the following topics, see [Security Dashboards](#):
  - DNS Top Tunneling Activity.
  - DNS Tunneling Traffic by Category.
  - Top Malware and DNS Tunneling Events by Client.
- **Syslog:** All DNS tunneling activities are logged to the syslog. You can view this log to identify specific activities related to DNS tunneling. For more information, see [Using a Syslog Server](#).

## Cisco ISE Integration

This section describes how to integrate Cisco ISE (Identity Services Engine) into the NIOS appliance to enhance identity management across devices and applications that are connected to your network routers and switches. It provides guidelines about how to subscribe and publish contextual data, and add notification rules. By integrating Cisco ISE, you can gain greater visibility in user and device information, thereby enhances security across your network resources. It includes the following topics:

- [Integrating Cisco ISE into NIOS](#)
- [Configuring Cisco ISE on NIOS](#)
- [Publishing Data](#)

## Integrating Cisco ISE into NIOS

With the rapid growth of BYOD (Bring Your Own Device) trend, the complexity of securing network resources has become more challenging. To ensure data privacy and security of all network resources against threats, Infoblox introduces the **Ecosystem** feature that allows you to expand the visibility of networks, users, and devices. Using this feature improves overall IT operations by sharing information between network and security teams. Integrating Cisco ISE server into NIOS enables NIOS and Cisco ISE to exchange valuable network, user, device, and security-event information, enriching both Infoblox DDI and Cisco ISE data. Cisco ISE is a centralized security solution (Network Access Control) that automates and enforces context-aware security access to network resources. This feature ensures that only the authorized users from legitimate devices get access to the services they need.



### Note

- Cisco ISE does not support IPv6 addresses.
- Cisco ISE version 3.0 is only supported in NIOS version 8.6.2.
- If you want to use Cisco ISE version 3.1 (pxGrid 2.0), Infoblox recommends that you integrate using the Outbound endpoint. For more information about outbound endpoints, see [Configuring Outbound Endpoints](#).

When you configure a Cisco ISE, you can do the following:

- **Subscribe to contextual data:** NIOS acts as a client to the Cisco ISE and collects information about the subscribed data types. You can configure extensible attributes without restricting them to specific object types, and then map these extensible attributes to Cisco ISE data to collect additional information. You can view subscribed information collected from the Cisco ISE in the appropriate tabs (**IPAM**, **IP Map** panel, and **Network Users**) of the Infoblox GUI. For information about how to subscribe to contextual data, see [Configuring Cisco ISE on NIOS](#). You can also monitor subscription data using the **Subscription** report. For information, see [Subscription Data](#).
- **Publish contextual data** - You can publish contextual data from NIOS to specific Cisco ISE based on the conditions and criteria specified in the notification rules. To publish RPZ and threat protection notifications, you must first set up an external syslog server, as described in [Specifying Syslog Server for Notifications](#). For more information about configuring notification rules, see [Configuring Notification Rules](#). You can monitor published data using the **Publish Data** report through the Reporting and Analytics feature. For information, see [Publish data](#).

## Supported Integrations

NIOS supports the integration of Cisco ISE versions 2.6 and 2.7, and 3.0.

Infoblox recommends that you use the Outbound endpoint for Cisco ISE integrations using Cisco pxGrid 2.0. For more information about outbound endpoints, see [Configuring Outbound Endpoints](#).

## Administrative Permissions

By default, only superusers can add, edit, and delete Cisco ISEs. Limited-access admin groups can access Cisco ISEs only if their administrative permissions are defined. For information about administrative permissions, see [About Administrative Permissions](#).

## Prerequisites to Integrate Cisco ISE with NIOS

Do the following before you begin using this feature on NIOS:

- You must install the **Security Ecosystem** license to configure Cisco ISE. You might need the following licenses to configure notification rules for RPZ and threat protection event types:

License	Event Types
RPZ	DNS RPZ
Threat Protection	Security ADP
DNS, DHCP, and MSMGMT	IPAM
DNS and DHCP	DHCP Lease

For information about how to install licenses, see [Managing Licenses](#).

- Cisco ISE uses SSL certificates as the method of authentication. You must upload the client certificate and client key when configuring the Cisco ISE server. You can include both client certificate and key in a single file and then upload. For information, see [Generating Certificates](#).

Note to make sure to use the host name of the Grid member that is selected as the subscribing member. The host name of the subscribing member must match with the Common Name that you mention while generating the certificate.

- For the bulk download certificate, download the server certificate from the monitoring node. If the admin node and monitoring node are on one node, then download the certificate from the admin node. Log into Cisco ISE and download the default self-signed server certificate (**Administration -> System -> Certificates -> Export**).
- For the CA certificate, download the CA certificate from the admin node or the self-signed certificate (**Administration -> System -> Certificates -> Export**).
- Register NIOS as a client on the Cisco ISE. You must enable the **Auto-Registration** option on the Cisco ISE: From the **Administration** menu -> click **pxGrid Services**, and then click **Enable Auto-Registration**. For more information, refer to Cisco ISE documentation. When you register NIOS successfully, you can view `infoblox_client_subscribe_xxxx` and `infoblox_client_publish_xxxx`, where `xxxx` is a number generated based on the IP of the subscribing member on the Cisco ISE. If auto-registration is not enabled, approve the pxGrid client after registration. If you change the certificates, Cisco ISE may not register the client successfully. In this case, delete the related pxGrid client from the Cisco ISE server, which is automatically created again.
- Enable the Identity Mapping feature on the NIOS appliance:
  - From the **Grid** tab, select the **Grid Manager** tab -> **Grid Properties -> Edit** from the Toolbar.
  - In the *Grid Properties Editor*, select the **General** tab -> **Advanced** tab, select the **Enable network users feature** checkbox.
- To publish data:
  - To publish dynamic data, such as DHCP lease and IPAM information, make sure that you approve **Infoblox\_DHCP** and **Infoblox\_IPAM** on the Cisco ISE, and then configure notification rules as described in [Configuring Notification Rules](#).
  - To publish RPZ and threat protection notifications to the Cisco ISE server, you must first set up an external syslog server and then configure notification rules, as follows:

- i. Configure an external syslog server that listens on port 2000, as described in [Specifying Syslog Server for Notifications](#).
- ii. Set up notification rules, as described in [Configuring Notification Rules](#).



#### Note

Refer to Cisco ISE documentation for information about how to perform auto-registration, creating authorized groups, and approving dynamic topics.

## Limitations of Integrating Cisco ISE with NIOS

Integrating Cisco ISE with NIOS has the following limitations:

- You can publish IPAM data only from the Grid Master that is a subscribing member. A subscribing member is a Grid member that you want to subscribe as the client on the Cisco ISE. For more information, see [Publishing Data](#).
- Only the subscribing member can publish its data to Cisco pxGrid.
- If the Grid Master is the subscribing member and you promote a Grid Master candidate to the Grid Master, then you have to create a client certificate for the promoted Grid Master.

## Generating Certificates

To generate a self-signed key and certificate:

1. `openssl genrsa -out self1.key 4096`
2. `openssl req -new -key self1.key -out self1.csr`
3. `openssl req -x509 -days 365 -key self1.key -in self1.csr -out self1.cer`

For CSR request:

Country Name (2 letter code) [XX]: <Country Name>, for example: US

State or Province Name (full name) []: <State Name>, for example: CA

Locality Name (eg, city) [Default City]:<City Name>, for example: SC

Organization Name (eg, company) [Default Company Ltd]:<Company Name>, for example Infoblox

Organizational Unit Name (eg, section) []:<Organization Name>, for example: QA

Common Name (eg, your name or your server's hostname) []:<host name of the subscribing member>

Email Address []:

Enter the following 'extra' attributes to be sent with your certificate request:

A challenge password []:

Import the certificate generated in step 3 to Cisco ISE's trusted store. Select the **Trust for authentication within ISE** checkbox.


Export the self-signed ISE certificate of the ISE server (under System -> Certificates). Make sure to select the **pxGrid: Use certificate for the pxGrid Controller** checkbox before exporting it.

You can call this as isemnt.cert

Wait for ISE services to restart. It may take a few minutes.

## Configuring Cisco ISE on NIOS

You can configure the supported versions of Cisco ISE servers on the NIOS appliance. You can subscribe for identity information that you wish to collect from the Cisco ISE, such as user name, domain name, VLAN, session state, SSID, endpoint profile, and security group. You can also add extensible attributes without restricting it to specific object types, and map these extensible attributes with the Cisco ISE field types to collect additional information. Note that you can subscribe to only one Cisco ISE per member and each member can subscribe to only one Cisco ISE. You can publish ADP and RPZ notifications, DHCP and IPAM information from NIOS to Cisco ISEs based on the notification rules that you have configured. You can view the subscribed information from the IPAM tab and the IP Map panel. Make sure that you synchronize time between the managing member and Cisco ISE.

 The procedures in the sections below is for configuring Cisco pxGrid 1.0. For instructions on configuring Cisco pxGrid 2.0, see [Configuring Outbound Endpoints](#). Infoblox recommends that you integrate Cisco ISE through the outbound endpoint.

### Configuring Cisco ISE Servers

You can configure a Cisco server either by using the **Ecosystem -> Cisco ISE Endpoint** tab or by using the **Ecosystem -> Outbound Endpoint -> Add -> Add Cisco ISE Endpoint** option. This section describes how to configure a Cisco server using the **Ecosystem -> Cisco ISE Endpoint** tab. For information about using the **Add Cisco ISE Endpoint** option, see [Configuring Outbound Endpoints](#).

To configure a Cisco ISE server:

1. From the **Grid** tab, select the **Ecosystem** tab -> **Cisco ISE Endpoint** tab, and then click the Add icon.  
or  
From the **Grid** tab, select the **Ecosystem** tab, and click **Add Cisco ISE** from the Toolbar.
2. In the *Add Cisco ISE* wizard, complete the following.
  - **Server Address**: Enter the IP address of the Cisco ISE.
  - **Version**: Select the version of the Cisco ISE.
  - **Subscribing Member**: Click **Select** to select a Grid member that you want to subscribe as the client on the Cisco ISE. In the *Member Selector* dialog box, select a Grid member from the list. This member interacts with the Cisco ISE to obtain contextual information for the subscribed data types.
  - **Network View**: This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the network.
  - **Client Certificate**: Click **Select** to upload the client certificate. In the *Upload* dialog box, click **Select** to navigate to the certificate, and then click **Upload**.
  - **Bulk Download Certificate**: Click **Select** to download the server certificate from the monitoring node or self-signed certificate.
  - **Manage Certificates**: Click **CA Certificates** to upload the self-signed certificate or CA certificate. In the *CA Certificates* dialog box, click the Add icon, and then navigate to the certificate to upload it.
  - **Test Credentials**: Click this to validate the Cisco ISE configuration before proceeding. When you click **Test Credentials**, the appliance validates the certificates.
  - **Comment**: Enter additional information about the configuration.
  - **Disabled**: Select this if you want to save the configuration but do not want to use it yet. You can clear this checkbox when you are ready to use this Cisco ISE.
3. Click **Next** to specify the data types that you are interested to obtain from the Cisco ISE. The Cisco ISE shares information only for the subscribed data types. Complete the following to specify data types you want to collect from the Cisco ISE:
  - **Subscription Settings**: There are predefined data types in the **Available Data Type** table you can subscribe. Use the arrows to move data types from the **Available Data Type** table to the **Selected Data**



**Type** table and vice versa. The appliance receives information for all data types in the **Selected Data Type** table.

- **Map other data types to Extensible Attributes:** You can create extensible attributes and map these extensible attributes to receive additional Cisco ISE data values, such as IP address, MAC, NAS IP Address, NAS Port ID, EPS Status, Posture Status, Posture Timestamp, Endpoint Profile Name, Account Session ID, and Audit Session ID. Click the Add icon and map a Cisco ISE data type to an extensible attribute. You can also select a row and click the Delete icon to delete it.
4. Click Next to add extensible attributes to the Cisco ISE. For information, see [Managing Extensible Attributes](#).
  5. Save the configuration.

## Modifying Cisco ISE Configurations

You can select data types that need to be published from NIOS to Cisco ISE after you have configured the Cisco ISE. You can modify the Cisco ISE configurations, as follows:

1. From the **Grid** tab, select the **Ecosystem** tab -> **Cisco** tab, click the **Action** icon next to the server name and select **Edit** from the menu.
2. The *Cisco ISE Server* editor provides the following tabs from which you can modify data:
  - **General:** You can modify data in this tab as described in [Configuring Cisco ISE on NIOS](#).
  - **Subscription:** You can edit data types that you have subscribed. You can use the arrows to move data types from the **Available Data Type** table to the **Selected Data Type** table and vice versa. The appliance receives information for all data types in the **Selected Data Type** table and extensible attributes that are configured.
  - **Publication:** To publish dynamic data from NIOS, you must first configure notification rules, as described in [Configuring Notification Rules](#). You can add data types that you want to publish to Cisco ISE server by using the arrows to move data types from the **Available** table to the **Selected** table and vice versa. The appliance publishes information only for the data types that are added in the **Selected** table.
  - **Extensible Attributes:** You can add, modify, and delete extensible attributes that are associated with the Cisco ISE server. For information, see [Managing Extensible Attributes](#).
3. Save the changes.

## Overriding Subscription Settings

You can override subscription settings and mapped extensible attributes at the network container, network, and DHCP range levels. By default, networks inherit subscription settings from those configured while adding the Cisco server. You can override these settings and subscribe new values at the DHCP range, network container, or network level. A network inherits subscription settings from its parent object. If you override the values at the network container level, then the network inherits the network container values. Otherwise, the network continues to inherit the values configured from the Cisco ISE. A shared network without a parent network container continues to inherit settings from the Cisco ISE.

To override an inherited value, click **Override** next to it and complete the appropriate fields. When you click **Override**, the appliance displays the value inherited from its parent object (if any).

To override subscription settings and mapped extensible attributes:

1. **Network Level:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* checkbox, and then click the Edit icon.  
**Network Container:** From the **Data Management** tab, select the **IPAM** tab -> *network\_container* checkbox, and then click the Edit icon.  
**DHCP Range Level:** From the **Data Management** tab, select the **DHCP** tab -> **Networks** tab -> **Networks** -> *network* -> *addr\_range* checkbox, and then click the Edit icon.
2. In the *Network or Range* editor, click **Toggle Advanced Mode** if the editor is in basic mode, and then click the Cisco ISE tab.
3. Save the configuration and click **Restart** if it appears at the top of the screen.

## Viewing Identity Mapping Information

To view user information, you must first enable identity mapping feature at the Grid level. For information about enabling Identity Mapping feature, see [Enabling Identity Mapping](#).



 **Note**

You do not need an **MSManagement** license to enable the identity mapping feature.

You can view user information in the **Network Users** tab. For more information, see [Viewing Identity Mapping Information](#).

### Deleting Cisco ISE Servers

When you delete a Cisco ISE, the appliance moves it to the Recycle Bin, if enabled. You can later restore it if needed. To delete a Cisco ISE server:

1. From the **Grid** tab, select the **Ecosystem** tab > **Cisco** tab -> *Cisco ISE server* checkbox, and then click the Delete icon.
2. In the *Delete Confirmation* dialog box, click **Yes** to delete the Cisco ISE server.

### Publishing Data

To publish dynamic data, such as DHCP lease and IPAM information, make sure that you approve **Infoblox\_DHCP** and **Infoblox\_IPAM** on the Cisco ISE, and then configure notification rules as described in [Configuring Notification Rules](#) below.

To publish RPZ and threat protection notifications to the Cisco ISE server, you must first set up an external syslog server and then configure notification rules, as follows:

1. Configure an external syslog server that listens on port 2000, as described in [Specifying Syslog Server for Notifications](#) below.
2. Set up notification rules, as described in [Configuring Notification Rules](#) below.

### Specifying Syslog Server for Notifications

Before you can publish RPZ and threat protection notifications to the Cisco ISE, you must first configure the syslog server to which the appliance logs RPZ and threat protection events. The appliance generate notifications about these events and analyze the data before sending it to the Cisco ISE. When setting up the syslog server, ensure that you select **DNS RPZ** and **Threat Protection** logging categories so all events related to RPZ and threat protection hits are logged to the syslog.



**Note**

For Cisco ISE to take appropriate action to quarantine malicious IP addresses, ensure that the **EPSSstatus** (Endpoint Protection Status) in the Authorization Policy is set to "**Quarantine**." This is set by default.

To specify an external syslog server in NIOS, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab -> **Members** tab, and then click **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties* editor, select the **Monitoring** tab, and then follow the procedures described in [Specifying Syslog Servers](#) using the following settings:
  - **Address:** Enter the loopback address **127.0.0.1** so that the appliance sends notifications to itself.
  - **Transport:** Select **UDP**.
  - **Interface:** Select **LAN**. The appliance uses the LAN1 port to send syslog messages.
  - **Source:** Select **Any**. The appliance sends both internal and external syslog messages.
  - **Port:** Enter **2000** as the port number.
  - **Logging Category:** Select **DNS RPZ** and **Threat Protection**.
3. Save the configuration.

## Configuring Notification Rules

You can configure notification rules after you have configured Cisco ISE on the NIOS appliance. For information, see [Configuring Cisco ISE on NIOS](#). To publish data and notifications from NIOS to Cisco ISE, you must configure notification rules. You can create notification rules for the following event types: **DNS RPZ**, **Security ADP**, **IPAM**, and **DHCP Lease**. Note that the DNS RPZ and Security ADP event types are available only if you have installed **RPZ** and **Threat Protection** licenses in the Grid. Each notification rule specifies the target Cisco ISE, the Grid member on which you wish to run this rule, notification rule criteria, and the action to be taken for the matching events. NIOS publishes information, such as DHCP lease information, IPAM data, and quarantine events, when the triggered events matches the notification rule criteria. Note that the **DHCP Lease** and **IPAM** event types are available only for the Cisco ISE 2.0 and 2.2 target servers.



### Note

Quarantine events are published to the Cisco ISE whenever the first rule matches the trigger criteria and it ignores all other rules.

To add notification rules:

1. From the **Grid** tab, select the **Ecosystem** tab -> **Notification** tab, and then click the Add icon.  
Or  
From the **Grid** tab, select the **Ecosystem** tab, and click **Add Notification Rule** from the Toolbar.
2. In the *Add Notification* wizard, complete the following.
  - **Name**: Enter the name of the rule.
  - **Target**: Select the IP address of the target server on which you want to publish from NIOS. This field displays all the IP addresses of the Infoblox servers and the PT servers that you added.
  - **Comment**: Enter useful information about the notification rule.
  - **Disable**: Select this option to disable the notification rule.
3. Click **Next** and complete the following:
  - **Event**: The appliance displays the list of event types based on the licenses installed. The values in the drop-down list are:
    - **DNS RPZ**: Select this to create notification rules for the DNS RPZ events.
    - **Security ADP**: Select this to create notification rules for Security ADP threat events.
    - **IPAM Type**: Select this to send IPAM data. No notification rule is required for this event type.
    - **DHCP Leases**: Select this to create notification rules for DHCP Lease events. This is available for Cisco ISE 2.0 and 2.2 servers.
  - In the **Match the following rule** section, select filters, operators and values from the drop-down lists for the selected event type. You can use the + icon to construct nested expressions within an event category. See the table detail at the end of this procedure.
  - For **IPAM** and **DHCP Lease** events: In the *Notify the target* section, there are predefined data types in the **Available** table you can publish. Click **Override** and use the arrows to move data types from the **Available** table to the **Selected** table and vice versa. The appliance sends information for all data types that are added to the **Selected** table. If you do not override, the publication settings is inherited from those configured while adding the Cisco ISE server. Note that you can configure only one IPAM rule per Cisco ISE server.
  - **Action**: The action to be taken for various events. Displays **Quarantine the end host** for **DNS RPZ** and **Security ADP** events. The Quarantine the end host action and Notify target data action are published through the subscribing member. Only the subscribing member can publish data to the Cisco pxGrid node.
4. Click **Next** to select Grid members. You can apply this notification rule on specific Grid members or apply this notification rule on all the Grid members.
  - **Apply rule to relevant members**: Select this option to apply notification rule to all relevant Grid members.
  - **Select Member(s)**: Select this option to select a Grid member for applying the notification rule. If there are multiple members, the *Member Selector* dialog box is displayed, from which you can select a member. Click the required member name in the dialog box. You can also click **Clear** to clear the displayed member and select a new one.
5. Click **Save** to save the Cisco ISE configuration.

Event Type	Filters	Operators	Value
<b>DNS RPZ</b>	Query Name	equals, begins with, and ends with	Enter the value that you want your rule to match
	Rule Name	equals, begins with, and ends with	Enter the value that you want your rule to match
	Action Policy	equals	Log Only, None, Block No Data, Block No Such Domain, Passthru, Substitute Domain Name
	Source IP	equals, matches CIDR, matches range	Enter the value that you want your rule to match
<b>Security ADP</b>	Rule Severity	equals, equal to or more severe, equal to or less severe	Information, Major, Critical, Warning
	SID	contains, equals, begins with and ends with	Enter the value that you want your rule to match
	Rule Message	contains, equals, begins with and ends with	Enter the value that you want your rule to match
	Source IP	equals, matches CIDR, matches range	Enter the value that you want your rule to match
<b>DHCP Leases</b>	Lease State	equals	Started, Renewed, and Expired
	You can override your Publish settings configured for the Cisco ISE server.		

## Examples

The following illustrations show sample notification rules and how the information is displayed in Grid Manager and the Cisco ISE:

### Sample Notification Rule for RPZ Events

Add Notification Wizard > Step 2 of 4

It may take up to a minute to apply the new rules.

\*Event

DNS RPZ

Match the following rule:

Reset

Action Policy

equals

Passthru

Cancel

Previous

Next

Save & Close

### Matching DNS RPZ Events

The screenshot shows the 'default' Response Policy Zone configuration page. The breadcrumb trail is: Dashboards > Data Management > Smart Folders > Grid > Administration > DNS > Response Policy Zones. The page title is 'default'. There is a 'Quick Filter' set to 'None' and a 'Filter On' button. Below the filter is a table with columns: NAME, TYPE, SEVERITY, PRIMARY NAME S..., LAST UPDATED, COMMENT, MULTI-MASTER Z..., and SITE. The table contains one entry: 'test' with type 'Local' and severity 'Major'.

The screenshot shows the 'test' Response Policy Zone configuration page. The breadcrumb trail is: Dashboards > Data Management > Smart Folders > Grid > Administration > DNS > Response Policy Zones. The page title is 'test'. There is a 'Quick Filter' set to 'None' and a 'Filter On' button. Below the filter is a table with columns: NAME OR ADDRESS, POLICY, DATA, COMMENT, and SITE. The table contains one entry: 'google.com' with policy 'Passthru Domal...'.

## Ecosystem - Outbound Notifications

This section describes how you can use outbound templates to convert NIOS specific events into outbound notifications and send the notifications to REST (REpresentational State Transfer) enabled and DXL (Data Exchange Layer) endpoints. These endpoints use the REST APIs to send user and event data to external servers or to update other outbound settings. It also explains how to create and add RESTful API and DXL endpoints that use NIOS specific data to construct outbound notifications, and how to configure notification rules to trigger outbound notifications based on the parameters defined in the templates. It also includes sample templates for supported servers.

It includes the following topics:

- [Outbound Notification Overview](#)
- [Configuring Outbound Notifications](#)
- [About Outbound Templates](#)
- [Creating Session Management Templates](#)
- [Creating Action Templates](#)
- [Configuring Outbound Endpoints](#)
- [Configuring Notification Rules](#)
- [Sample Templates](#)
- [Configuring BloxOne Threat Defense Cloud Clients for Outbound](#)
- [Configuring Outbound Cloud Clients](#)

### Outbound Notification Overview

The Infoblox outbound API is a framework that is used to exchange both IPAM data (such as networks, network containers, hosts, leases) and DNS threat data with external interfaces. It sends object information and conversations to other REST APIs when an event triggers in NIOS. It is important that you receive notifications about the updates to the system. On the other hand, you may sometimes also need to identify and manage low-risk or accidental threats so the endpoint performance is not negatively affected. For example, if a user inadvertently browses to a faulty web site and you have configured RPZ rules to block this site, you may want to receive notifications and take certain actions so the user is not being blocked or quarantined. In addition, when the Infoblox appliance detects a new host or network, the detection might trigger a vulnerability scan by services such as Qualys and a scan for RPZ events configured in NIOS. In this scenario, you might want to configure conditions to capture these events so you can receive outbound notifications and perform appropriate actions to handle the situation.

For example, you can first configure RPZ rules to mitigate a malicious IP address, and then configure RESTful API and DXL endpoints to which you want to send the outbound notifications. When configuring your notification rules, you can match RPZ events that are initiated by the RPZ rules and apply the outbound template containing actions to mitigate the threat. The configuration rules then trigger outbound notifications, and the appliance sends the notifications to the configured endpoints and applies the configured actions to combat the offensive IP address.

Before you configure the appliance to send outbound notifications, there are a few limitations you might want to consider, as described in [Best Practices for Outbound Notifications](#), below. For detailed information about how to use the outbound notification feature, see [Configuring Outbound Notifications](#).

To enable, configure, and test outbound notifications you must first install the **Security Ecosystem** license in your Grid and do the following in the below mentioned order:

1. Ensure that the necessary services and features are configured. These include DHCP, RPZ, Threat Insight, ADP, Network Insight, and BloxOne Threat Defense Cloud.
2. Create necessary extensible attributes, if required. For more information, see [Managing Extensible Attributes](#).
3. Create or download login and logout templates from the Infoblox Community Site at <https://community.infoblox.com>. Next, add or upload the login and logout templates followed by the session template. Note that you can add a session template or download it from the Infoblox Community Site.
4. Download or create notification templates from the Infoblox Community Site at <https://community.infoblox.com>. Next, add or upload the notification templates.

5. Add an endpoint. You can either add REST API or DXL endpoints. For DXL endpoints, you must generate a NIOS client certificate, import DXL certificates and import or add list of DXL brokers. For more information, see [Configuring Outbound Endpoints](#).
6. Define notification rules. For more information, see [Configuring Notification Rules](#).

The outbound notification feature employs the following mechanism to enable and deliver event-driven messages to configured endpoints:

1. Accepts the configuration of events that you want to monitor (such as RPZ hits) and the configuration of endpoints to which you want to send outbound notifications.
2. Filters events for specific data sets or thresholds, such as RPZ hits for a specific domain within a specific time interval.
3. Matches the selected events and conditions defined in the templates to create outbound messages.
4. Sends outbound notifications to the configured endpoints.

Depending on the notification rules for RPZ and threat protection event types you want to configure on NIOS, you may need to install the applicable licenses. For information about other licensing requirements, see [Licensing Requirements](#) below.

 **Note**

To access online resources about this feature, including training videos and sample outbound templates for supported ecosystem partners, ensure that you visit the Infoblox Community Site at <https://community.infoblox.com>.

For debugging purposes, you can look at the syslog to see if the Outbound service has been started or stopped on specific members. You can also set the logging level to **Debug** to view all events in the log files, including deduplicated events. However, leaving the logging level at the **Debug** level could negatively affect your system performance. Therefore, Infoblox does not recommend leaving the logging level at **Debug**. For information about how to configure the severity level and deduplication, see [Configuring Outbound Endpoints](#).

## Licensing Requirements

You must install the **Security Ecosystem** license to enable outbound API notifications. After you install the Security Ecosystem license, you can configure REST and DXL endpoints. If you do not have this license installed, the outbound notification feature is disabled. You might also need the following licenses to configure notification rules for certain event types:

*Licenses required for certain event types*

License	Event Types
RPZ	DNS RPZ
DNS and DHCP	DHCP Lease
Threat Analytics	DNS Tunneling
Advanced DNS Protection	Security ADP
RPZ and Security Ecosystem	BloxOne Threat Defense Cloud
Network Discovery	Object Change Discovery Data

For information about how to install licenses, see [Managing Licenses](#).

## Administrative Permissions

Only superusers can add, edit, and delete REST endpoints and notification rules by default. Limited-access admin groups can perform these tasks only if their administrative permissions are defined. For information about administrative permissions, see [About Administrative Permissions](#).

## Best Practices for Outbound Notifications

The following are some best practices and limitations you might want to consider while configuring outbound notifications:

- You can configure REST and DXL endpoints only on the Grid Master and Grid Master Candidate, but not on Grid members.
- During a scheduled full upgrade in the Grid, you cannot modify any configuration related to the outbound feature until all Grid members are upgraded.
- Outbound notification is not supported during an HA failover. Any events that are in transit during a failover might be lost.
- When you remove or disable a notification rule, no new events will be triggered. However, the appliance continues to process events that are already in queue.
- The buffer to hold events temporarily are limited and not configurable in this release. In very unlikely conditions, events may be dropped due to a full buffer. If events are dropped, summary information is logged to the syslog to indicate the type of events and the number that have been dropped. If this issue occurs continuously, contact Infoblox Technical Support.
- Events generated due to changes made by admin users do not support the Microsoft Management feature. The appliance does not generate events when there are changes done from the Microsoft servers. However, if you make changes that need to be synchronized to the Microsoft servers, the object change event is generated before the changes are synchronized with the Microsoft servers.
- The Grid Master Candidate will continue to perform event enrichments and outbound API calls during and after a Grid Master promotion.
- If you disable the outbound notification feature or make changes to stop future notifications sent to an endpoint, all notifications that are currently in queue for this endpoint will stop immediately.
- The appliance uses rate limiting to control both data collection from Grid members and outbound notifications to external endpoints. It is possible for the appliance to drop events if its buffer is full or if there is a loss of connection between the Grid Master and the Grid members. Logs for these events are consolidated and logged to the syslog.
- The number of outbound notifications sent to external endpoints can be limited, depending on the requirements configured for the external servers. For example, some REST enabled servers only take 10 API calls per second. Some servers might put a user in suspended mode if the number of API calls sent to the user exceeds the limit. If necessary, you can adjust the rate limit criteria for API calls on the external servers.

## Configuring Outbound Notifications

A notification is an association between an event type, a template, and an endpoint. You must define the event type that triggers the notification, choose the template you want to use and an API endpoint with which the Grid must establish a connection.

Note the following before you configure outbound notifications:

- You must define an endpoint and create or upload outbound templates to the Grid before adding a notification.
- To send outbound notifications to an endpoint, you must configure notification rules first. A notification rule contains the target endpoint to which you want to send outbound notifications and the event type upon which you want to take action.
- The event type you select in a notification rule must match the event type defined in the template you want to use for that rule. Otherwise, the appliance returns an error.

To configure outbound notifications, complete the following:

1. Prepare outbound templates that you want to use for notification rules. For more information about API templates, see [About Outbound Templates](#). You can also reference sample API templates for supported ecosystem partners and modify them accordingly.
2. Upload outbound templates to the Grid so you can use them for corresponding notification rules. For more information, see in [Adding Outbound templates](#).
3. Configure endpoints to which you want to send outbound notifications. Infoblox supports the following endpoints: RESTful API and DXL. For more information, see [Configuring Outbound Endpoints](#).
4. Select event types and configure criteria for notification rules. For more information, see [Configuring Notification Rules](#).

**Note**

To access online resources about this feature, including training videos and sample API templates for supported ecosystem partners, visit the Infoblox Community Site at <https://community.infoblox.com>.

## About Outbound Templates

The appliance uses outbound templates to convert NIOS events into REST API, DXL, and Syslog messages. Outbound API templates control the integration and contains necessary steps to execute outbound notifications. Infoblox supports session management and action templates that are available in the Infoblox Community Site at <https://community.infoblox.com>. You can either modify these existing templates or create new ones through Grid Manager using the existing templates. You use supported variables in these templates to get respective events and define actions you want to take for those events. You can also define additional extensible attributes that are necessary for certain templates. The appliance uses the scripts in these templates to process and send outbound notifications to the endpoints.

The following are the outbound templates that you can create:

- **Session Management Template:** A session management template contains specific variables about an endpoint, such as the timeout value and rate limiting information. For more information, see [Creating Session Management Templates](#).
- **Action Template:** An action template defines the action(s) to be taken on the selected endpoint for the matching event type(s). It contains scripts the appliance uses to query respective event data from NIOS and to perform actions you want to take in response to the events. For more information, see [Creating Action Templates](#).

**Note**

When you add an outbound template or make changes to an existing one, it may take a few seconds to a few minutes until the changes are propagated to all the members.

You can also export the schema from the appliance and use it to create the outbound templates. For information about how to export template schema, see [Exporting Template Schema](#) below. Note that the exported schema is in the IETF JSON Schema format. For information about this schema format, see <https://tools.ietf.org/html/draft-zyp-json-schema-04>.

**Note**

To access online resources about this feature, including training videos and sample outbound templates for supported ecosystem partners, visit the Infoblox Community Site at <https://community.infoblox.com>.



## Adding Outbound Templates


You can use the supported templates that are available in the Infoblox Community Site at <https://community.infoblox.com>. You can download these templates make necessary changes and upload them to the Infoblox Grid.

Complete the following to add or upload an outbound template to the Infoblox Grid:

1. **Grid:** From the **Grid** tab, select the **Ecosystem** tab -> **Templates** tab, and then click the Add icon.  
or  
From the **Grid** tab, select the **Ecosystem** tab, and click **Add Template** from the Toolbar.  
**System:** From the **System** tab, select the **Ecosystem** tab -> **Templates** tab, and then click the Add icon.  
or  
From the **System** tab, select the **Ecosystem** tab, and click **Add Template** from the Toolbar.
2. In the *Add Template* wizard, complete the following:
  - Click **Select** to upload an outbound template. In the **Upload** dialog box, click **Select** and navigate to the template, and then click **Upload**. Select the **Overwrite the existing template** checkbox to overwrite an existing template.
3. Click **Add** to add an outbound template. To add a session management template, see [Creating Session Management Templates](#) or [Creating Action Templates](#) to create an action template.
4. Optionally, click **View Results** to open the *Syslog Preview* dialog and view all the syslog messages. For more information, see [Previewing Syslog Events](#).

## Modifying Outbound Templates

To modify an outbound template:

1. **Grid:** From the **Grid** tab, select the **Ecosystem** tab -> **Templates** tab -> *template* checkbox, and then click the Edit icon.  
or  
From the **Grid** tab, select the **Ecosystem** tab, select the template that you want to modify, click the Action icon  and select **Edit**.  
**System:** From the **System** tab, select the **Ecosystem** tab -> **Templates** tab -> *template* checkbox, and then click the Edit icon.  
or  
From the **System** tab, select the **Ecosystem** tab, select the template that you want to modify, click the Action icon and select **Edit**.
2. The *<Template Name>Template* editor contains the following tabs from which you can modify information:
  - **General:** You can modify the **Name** and **Comment** fields. All other fields are automatically propagated with available information, such as the template type, vendor type, event type, and action type. The **Name** field is optional for NIOS 8.3.0 and later releases, but note that it is mandatory for jump operations. When you do not specify a name for the template, an auto-generated value is set for this field. Example:  
`templatestep#N` or `functionstep#N` where N represents the step number. When you import a template, ensure that the goto functions are sent to named steps and not to automatically named steps.
  - **Contents:** This tab displays the content of the uploaded template file. You can modify the template contents and the appliance validates the content for proper JSON format when you save the configuration. For more information about the format of the templates, see [Creating Session Management Templates](#) and [Creating Action Templates](#).
3. Save the configuration.


## Viewing All Outbound Templates

To view the list of outbound templates:

1. **Grid:** From the **Grid** tab, select the **Ecosystem** tab, and then click the **Templates** tab.  
**System:** From the **System** tab, select the **Ecosystem** tab, and then click the **Templates** tab.
2. Grid/System Manager displays the following information:

- **Name:** The outbound template name.
- **Vendor Type:** The vendor type.
- **Event Type:** The event type specified in the template.
- **Template Type:** Displays the template type such as **Session Management** or **Event**.
- **Outbound Type:** The endpoint type. This can be REST, DXL, or Syslog.
- **Comment:** Comments that were entered for the outbound template.
- **Added On:** Displays the timestamp when the template was uploaded to the Grid Master in this format: YYYY-MM-DD HH:MM:SS, plus time zone.

You can do the following in this tab:

- You can select the Action icon  and do the following:
  - **Edit:** Select this to modify the outbound template information.
  - **Delete:** Select this to delete a template.
  - **Export:** Select this to export a template.
- Edit the outbound template information.
  - Select the outbound template, and then click the Edit icon.
- Delete an outbound template.
  - Select the template, and then click the Delete icon.
- Export the list of outbound templates.
  - Click the Export icon.
- Print the list of outbound templates.
  - Click the Print icon.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria:
  - In the filter section, click **Show Filter** and define filter criteria for the quick filter.
  - Click **Save** and complete the configuration In the Save Quick Filter dialog box.
 

The appliance adds the quick filter to the quick filter drop-down list in the panel. Note that global filters are prefixed with [G], local filters with [L], and system filters with [S].
- Sort the outbound templates in ascending or descending order by column.

## Deleting Outbound Templates

You can delete a template any time after you have created it.

### Note

Before you delete a template, ensure that it is not in use by a notification or an endpoint.

To delete a template:

1. From the **Grid/System** tab, select the **Ecosystem** tab, and then click the **Templates** tab.
2. Select the template that you want to delete, and click the Delete icon.
3. Click **Yes** in the *Delete Confirmation* dialog.

## Exporting Template Schema

The Export Template Schema feature allows you to export the session management and action template schema. You can use the exported schema to validate the templates before uploading them. The exported schema is in IETF JSON Schema format. It is a JSON file that is used to define and validate outbound templates that you want to upload to the appliance. If you want to create your own templates, you must follow this schema format. For information about this format, see <https://tools.ietf.org/html/draft-zyp-json-schema-04>.

Complete the following to export the outbound template schema:

1. From the **Grid/System** tab, select the **Ecosystem** tab -> select the **Templates** tab.
2. Expand the Toolbar, click **Export Template Schema** -> select **Action Template Schema** to export the action template schema or select **Session Template Schema** to export the session management template schema.

- If there is only one template version, the appliance downloads the schema to your local system. If there are multiple template versions, the appliance displays the *Export Action Template Schema* or *Export Session Template Schema* dialog. In the dialog, select the outbound template and the template version that you want to export. The template schema is downloaded to your local system.

## Creating Session Management Templates

You can use a session management template to specify settings that will be applied to an endpoint. You can define settings such as the timeout value after which the outbound requests are aborted. This template can also contain additional child templates that can be referenced by the configuration. However, you cannot reference other templates in the system from the session management template.

Once you upload a session management template to the Grid, the configuration in the template automatically applies to the specified endpoint, if the connection to the endpoint has been established and if the template is assigned to the endpoint.

The table below lists the supported variables you can use in a session management template.

### Note

Changes made to the variables of the template through Grid Manager override the variables in the template. Although the latest template version is 3.0, template versions 2.0 and 1.0 continue to work with their respective syntax. Session management templates do not support instance variables. To access online resources about this feature, including training videos and sample outbound templates for supported ecosystem partners, visit the Infoblox Community Site at <https://community.infoblox.com>.

The following are the session management templates that you can create:

## Session Management Template Variables

### *Session Management Template Variables*

Variable	Type	Mandatory	Description
version	Must be 5.0	Yes	The version number of the template. Note that 5.0 is the latest version. The appliance still fully supports the previous schema version.
type	Must be REST_ENDPOINT for REST API endpoints Must be DXL_ENDPOINT for DXL endpoints Must be SYSLOG_EVENT for Syslog endpoints	Yes	The template type.
name	String	Yes	The template name.
vendor_identifier	String	Yes	The vendor identifier for an endpoint.
comment	String	No	Additional information.

Variable	Type	Mandatory	Description
path	String	No	Path string to append to what the user enters in the GUI.
override_path	Boolean	No	If present, the path above replaces the path the user enters in the GUI.
timeout	Integer	No	The timeout value in seconds. The minimum is 1 and the maximum is 3600. The default is 30.
keepalive	Boolean	No	The value can be <b>True</b> or <b>False</b> . The default is <b>False</b> .
keepalive_timeout	Integer	No	The keepalive timeout value in seconds. The connection is closed after this timeout. The minimum is 1 and the maximum is 300. The default is 5.
dxl_keepalive_timeout	Integer	No	The keepalive timeout value in seconds with a connected DXL broker. This controls the rate at which the client will send ping messages to the broker if there is no exchange of other messages. The minimum is 1 and the maximum is 3600. The default is 30. <b>Note:</b> This field is applicable only for the DXL templates.
dxl_topic	String	No	The DXL topic that is used to send data by DXL. <b>Note:</b> This variable is applicable only for the DXL templates
retry	Integer	No	The number of attempts to try to connect to the endpoint before considering the connection a failure (this covers only timeout/network errors). The default value is 2.
retry_template	Integer	No	The number of attempts the appliance retries the full template if it returns a failure (this covers a template returning anything but 200). The default value is 0 (no retry).
rate_limit	Float	No	The maximum number of messages (per second) that are sent to an endpoint. The default value is 0 (no rate limiting). You can enter a value less than 1.0 to have lower rate limits. For example, if you set the value to 0.5, the appliance sends less than 1 message in every 2 seconds.
rate_limit_requests	Integer	No	Use this along with <code>rate_limit_interval</code> to define the absolute number of requests for rate limiting. For example, if you set this to 10 and <code>rate_limit_interval</code> to 5000 milliseconds, no further requests are sent until the number of requests falls below 10 in 5000 milliseconds. So if 3 requests are sent at 0 second and 7 requests are sent at the 3rd second, no requests will be sent until the 5th second, where 3 requests are allowed for another 3 seconds. The default is 0 = no rate limiting. You can set this to an integer between 0 and 26784000000.
rate_limit_interval	Integer (in milliseconds)	No	Use this along with <code>rate_limit_requests</code> to define the rate limiting interval . The default is 0 = no rate limiting. You can set this to an integer between 0 and 2678400000 milliseconds.

Variable	Type	Mandatory	Description
rate_limit_reset	Integer (in epoch time)	No	<p>Use this together with <code>rate_limit_requests</code> and <code>rate_limit_interval</code> to reset the epoch time for rate limiting. The default is 0 = not resetting the epoch time for rate limiting. You can set this to an integer between 0 and 4102437600 milliseconds.</p> <p>Example: If <code>rate_limit_requests = 10</code>, <code>rate_limit_interval = 5</code> seconds, and <code>rate_limit_reset = 3 p.m. on 10 January 2017</code> (set in epoch time), 10 requests can be sent within 5 seconds after the reset time.</p> <p>For instance, 3 messages are sent at the 1st second and 7 messages sent at the 3rd second and no more messages can be sent in the next 2 seconds, which is within 5 seconds after 3 p.m. on 10 January 2017.</p>
endpoint_variables	List of VARIABLE structs	No	If specified these will be accessible via the S: namespace in templates.
inactivity_interval	Integer (in milliseconds)	No	A logout request is sent after this time interval, provided no other requests have occurred during this time period. The default is 0 = "disable." You can set this to an integer between 0 and 2 <sup>63</sup> -1 milliseconds.
login_template	String	No	The template that requires user login if there is no active session currently running for the endpoint. Only templates with <code>event_type</code> set to SESSION are supported for this. You must include the template name in the string. Ensure that you upload the login template if you plan to use this variable in the session management template.
logout_any_condition	Boolean	No	Specifies whether a logout request is sent, depending on whether <code>logout_status_code</code> or <code>logout_regexp</code> are set. The value can be <b>True</b> or <b>False</b> . The default is <b>False</b> . If this is set to <b>True</b> , a logout request is triggered if either <code>logout_status_code</code> or <code>logout_regexp</code> is set. If this is set to <b>False</b> , both <code>logout_status_code</code> and <code>logout_regexp</code> must be set.
logout_regexp	String	No	Any response returned by the endpoint that matches the regular expression in this field will trigger a logout request.
logout_status_code	Integer	No	The HTTPS response status code used for the provisional response will trigger a logout request. The default is 0 = "disable." Valid values are between 0 and 599.
logout_template	String	No	The template that is being executed after steps are executed or after various session duration constraints are met. You must include the template name in the string. Only templates with <code>event_type</code> set to SESSION are supported for this. Ensure that you upload the logout template if you plan to use this variable in the session management template.
requests_per_session	Integer	No	The number of requests that are sent before a logout request is sent. The default is 0 = "disable." You can set this to an integer between 0 and 2 <sup>63</sup> -1.

Variable	Type	Mandatory	Description
logout_only_at_template_end	Boolean	No	The value can be <b>True</b> or <b>False</b> . The default is <b>False</b> . If this is set to <b>True</b> , a logout request is sent only after the execution of the template has been completed to ensure that all template requests pertaining to a session are executed if that is required by the API. Therefore, it is possible to set <code>requests_per_session</code> to 1 and have each template call executed in its own session regardless of the number of requests it contains.
session_cooldown	Integer (in milliseconds)	No	When this is set, a login request will be sent only after the amount of time set for this has been elapsed after a logout. The default is 0 = "disable." You can set this to an integer between 0 and 2 <sup>63</sup> -1 milliseconds.
session_duration	Integer (in milliseconds)	No	The time interval after which a logout request is sent unconditionally after a login. For example, if you set this to 150 milliseconds, after a login, a logout will be sent after 150 milliseconds whether there is current traffic or not. The default is 0 = "disable." You can set this to an integer between 0 and 2 <sup>63</sup> -1 milliseconds.
wapi_rate_limit	Float	No	Independent of the endpoint rate limiting, use this for persistent WAPI requests only. The maximum number of messages (per second) that are sent to an endpoint. The default value is 0 (no rate limiting). You can enter a value less than 1.0 to have lower rate limits. For example, if you set the value to 0.5, the appliance sends less than 1 message in every 2 seconds. The default is 0 = "disable." You can set this to an integer between 0 and 100000.
wapi_rate_limit_requests	Integer	No	Independent of the endpoint rate limiting, use this for persistent WAPI requests only. Use this along with the <code>wapi_rate_limit_interval</code> field to define the absolute number of requests for rate limiting. For example, if you set this to 10 and <code>rate_limit_interval</code> to 5000 milliseconds, no further requests are sent until the number of requests falls below 10 in 5000 milliseconds. So if 3 requests were sent at 0 second and 7 requests were sent at the 3rd second, no requests will be sent until the 5th second, where 3 requests are allowed for another 3 seconds. The default is 0 = "disable." You can set this to an integer between 0 and 26784000000.
wapi_rate_limit_interval	Integer (in milliseconds)	No	Independent of the endpoint rate limiting, use this for persistent WAPI requests only. Use this along with the <code>wapi_rate_limit_requests</code> field to define the rate limiting interval. The default is 0 = no rate limiting. The default is 0 = "disable." You can set this to an integer between 0 and 2678400000 milliseconds.
wapi_rate_limit_reset	Integer (epoch time)	No	Independent of the endpoint rate limiting, use this for persistent WAPI requests only. Use this to reset the epoch time for rate limiting. The default is 0 = not resetting the epoch time for rate limiting. You can set this to an integer between 0 and 4102437600 milliseconds.

## Creating Action Templates

The purpose of an action template is to convert an event into one or more RESTful API, DXL, and Syslog messages that are sent from the NIOS appliance to the endpoint configured in the notification rule. An action template consists of a series of statements that are interpreted into specific actions. When creating an action template, ensure that it consists of an initial section with some general template settings, followed by one or more steps that are executed in sequence.

Steps are constructed for sending messages to the endpoint and for receiving responses. It can also perform specific operations on template variables.



**Note**

The steps in an action template are executed sequentially. Some constructs enable steps to be skipped by jumping forward in the list of steps; jumping backwards is not supported.

Consider the following guidelines while composing action templates:

- Template error handling is active only if the severity level for logging is set to **Debug**, otherwise error handling is disabled and the server continues to execute a template even if the template tries to access nonexistent variables or perform invalid operations such as trying to increment a STRING variable. For information about setting the severity level for logging, see [Configuring Outbound Notifications](#). If you have disabled template error handling, then accessing any nonexistent variables will return an empty string and invalid operations are not executed.
- Matching a regular expression is performed un-anchored. If anchoring is required, you must add the characters ^ and/or \$ to the respective regular expression.
- For outbound notifications, only template instances are considered. Template instances are constructed from action templates as well as template instance variables in the template. You may configure these variables through Grid Manager when a specific template is associated to an event.
- You can use functions that contain the list of steps to be executed. The **FUNCTION** operation executes the steps in the function.
- Note that the function steps use step execution limit. Step counters for functions are removed when you return to template steps from function steps.
- You must specify the list of functions in the event template and these functions can be used only by the event template.
- Jumping between functions and template steps is not supported.
- Infoblox supports unused functions and allows execution of a function that is located within a function, including itself.
- You cannot upload or update a template if the function does not exist.

An action template consists of the variables and elements listed in the [Action template Variables](#) table below.



**Note**

When "yes" is indicated in the "Sub" column for a variable, it means that variable substitution (where it is possible to have \${...} variables as part of the value, and have them substituted when the template is executed) is supported for that variable, a \* means that the substitution is supported for some of the fields of the struct.

## Action Template Variables

### *Action template Variables*

Variable	Type	Mandatory	Sub	Description
version	Must be 5.0	No	No	The version number of the template. Note that 5.0 is the latest version. The appliance still fully supports the previous schema version.
name	String	No	No	The template name.

Variable	Type	Mandatory	Sub	Description
type	Must be REST_EVENT for REST API endpoints Must be DXL_EVENT for DXL endpoints Must be SYSLOG_EVENT for Syslog endpoints	No	No	The template type.
vendor_identifier	String	No	No	The vendor identifier for an endpoint.
event_type	List of ENUM (except for SESSION)	Yes	No	Available event types: RPZ, LEASE, TUNNEL, ADP, DXL, NETWORK_IPV4, NETWORK_IPV6, RANGE_IPV4, RANGE_IPV6, FIXED_ADDRESS_IPV4, FIXED_ADDRESS_IPV6, HOST_ADDRESS_IPV4, HOST_ADDRESS_IPV6, DISCOVERY_DATA, SCHEDULE, DNS_RECORD, DNS_ZONE, SESSION, SYSLOG_SEND_EVENT, DXL_SEND_REQUEST  Note that SESSION is used for the login and logout events for the session management templates. For information about supported variables for the session management template, see <a href="#">Session Management Template Variables</a> .
action_type	String	No	No	The action type.
comment	String	No	No	Additional information.
content_type	String	No	No	The content type for the whole template. If not specified, it is set as "application/json". It can be specified either in the template, inside the transport, or both. If both are specified, then the content type specified in the template takes precedence.
quoting	ENUM	No	No	Sets the default serialization for template variables. The valid value is one of the following: JSON, XML, XMLA, ASIS, or URL. If not specified, this is set to JSON. To use XMLA for serialization, you must specify a valid quoting for variables. Infoblox strongly recommends that you use XMLA as the quoting option (as opposed to XML) when you create new action templates. New operations such as PUSH, POP, SHIFT and others are not officially supported for XML parsed data.
headers	Dictionary of key/value pairs	No	Yes	If specified, it is sent for every request of the template, in addition to any other H: namespace variables.
instance_variables	List of VARIABLE structs	No	No	It must be specified in the GUI to create a template instance.
template_variables	List of VARIABLE structs	No	No	template_variables is used when there are common settings in various steps. They are specified for maintainability and they are also available in L: namespace. Note that the template variables are evaluated immediately so that they can be used as reference values in other areas of the template.
transport	TRANSPORT struct	No	No	It is the default for steps.



Variable	Type	Mandatory	Sub	Description
				Steps are executed in sequence. You cannot execute them
steps	List of STEP structs	Yes		Steps are executed in sequence. You cannot execute them in a backward manner.
step_execution_limit	Integer			The maximum number of times a step can be executed. Steps can be executed more than once if using certain template constructs. To avoid possible endless loops, you can limit an individual step to be executed up to a certain number of times. If a step is executed again after the limit is reached, the execution will be interrupted and a failure is logged. The default is 10. The maximum number is 1000.

## STEP Struct

Each step can perform a different operation, such as **SLEEP**, **CONDITION**, **NOP**, **GET**, **POST**, **PATCH**, **DELETE**, **DXL\_EVENT\_SEND**, or **PUT**. The STEP Structs table below, lists the step variables.

### Note

The name for each step must be unique.

- **SLEEP**: Steps with a SLEEP operation will pause the execution for the number of seconds specified in the timeout parameter.
- **NOP**: Steps with a NOP operation will only parse the text specified in body/body\_list without sending it to the endpoint. This can be useful for executing operations on variables to prepare data for subsequent steps.
- **CONDITION**: Steps with a CONDITION operation will cause a condition specified in statements to be matched. See **CONDITION Structs** table below, for information about condition variables and **STATEMENT Structs** table below, for statement variables.

If the condition matches, the execution of the template will do the following:

- Stop without errors if the stop field is present,
- Stop with an errors if the error field is present,
- Jump forward to the specified step if the next field is present,
- and/or evaluate the text in eval if the field is present.

If the condition does not match, the text in the else\_eval field will be evaluated instead, if the field is present.

- **GET**, **POST**, **PATCH**, **DELETE**, and **PUT**: These steps will result in an endpoint communication. The request will be sent to an URI composed by the URI configured in the endpoint, plus the path configured in the session management template if present (or replaced by it, depending on its override setting), plus the path configured in the event template if present (same override considerations), plus the path configured in the individual step if present (same override considerations).  
After variable substitution, any data present in the H: namespace will be sent as headers, any parameters listed will be sent as URI parameters, and any data in body/body\_list will be sent as the body of the request. Note also that each step can override the endpoint timeout. This is useful when a certain operation is known to require a longer execution time.  
If the result variable is set in the step, the reply from the server will be evaluated and compared against the specified codes and/or REGEX, and the operation specified in the result step will be executed.  
If parse is set in the step, the result from the server will also be parsed and be made available in the P: namespace as described above.

### Note

Set the "wapi" field to send WAPI requests to the Grid Master using available event data. For example, you can add or modify extensible attributes of a NIOS object at the time when the object is being synchronized. If you include the "wapi" field in a step, you must enable WAPI integration by entering the WAPI login username and password while configuring the endpoint. Otherwise, the WAPI step will fail due to an authorization error. For information about how to configure endpoints, see [Configuring Outbound Notifications](#).

### STEP Structs

Variable	Type	Mandatory	Sub	Description
name	String	Yes	No	Used to refer to the steps used in the execution of the template.
operation	ENUM	Yes	No	The valid value can be one of the following: GET, POST, DELETE, PATCH, PUT, SLEEP, CONDITION, NOP, VARIABLEOP, DXL_EVENT_SEND, SERIALIZE, or FUNCTION. If you specify SLEEP, only timeout is supported, where timeout is the sleep length in seconds. If you specify NOP, only variable operations are performed, and only body/body_list is supported. If you specify VARIABLEOP, you must use the VARIABLE struct within your steps that are executed in sequence. See VARIABLE Struct below, for the supported fields. If you specify <b>SERIALIZE</b> , see SERIALIZE Struct below, for the supported variables.
function_name	String	Yes		Function name. Applicable only if the operation is set to <b>FUNCTION</b> .
condition	CONDITION Struct	No		Applicable only if the operation is set to <b>CONDITION</b> .
timeout	String	No	Yes	If specified, overrides the endpoint configuration value (it is useful if the template is slow during execution). Note that the timeout value is invalid for NOP. Since the timeout variable is a string, you can substitute the variable in individual steps.
transport	TRANSPORT Struct	No	Path only	
result	List of RESULT Structs	No	No	If not present, you can assume 200, everything else is a failure. If not specified, the steps are executed sequentially. This is not valid for SLEEP, NOP, or CONDITION variables.
parse	ENUM	No	No	If specified, the output of the server will be parsed. The valid value is one of the following: JSON, REGEXLINE, REGEXMULTILINE, REGEX, REGEXSPLIT, XMLA, or XML. Infoblox recommends that you use <b>XMLA</b> instead of <b>XML</b> for parsing. Ensure that you see Result Parsing below for details.
parse_regex	String	No	No	You can set one of the following: REGEX, REGEXSPLIT, REGEXLINE, or REGEXMULTILINE

Variable	Type	Mandatory	Sub	Description
parameters	List of PARAMETERS	No	Value only	These are URI parameters.
headers	Dictionary of name/value pairs	No	Yes	This is sent as HTTP headers. The name space substitution is supported only for value. Note that assigning to the H: name space also sends headers.
override_headers	Boolean	No	No	If specified, only these headers and H: name space headers are sent instead of template headers.
body	String	No	Yes	This is applicable only for POST, PATCH, DXL_EVENT_SEND, and PUT requests as well as NOP and FUNCTION operations. It will be sent as the body of the request. Note that name space substitution is supported.
body_list	List of strings	No	Yes	This is an alternative to the body. If specified, the strings in the list will be joined before sending it. Any leading or trailing whitespace is removed. This is applicable to FUNCTION operation.
no_connection_debug	Boolean	No	No	The valid value is True or False. If this is set to True even if the endpoint is set to a Debug level logging, only the body, headers, and cookies for the corresponding step will NOT be output to the debug log. Only explicit DEBUG calls will be displayed. This is generally used in login templates to avoid usernames and passwords from being logged to the log files in plain text.
variable_ops	List of VARIABLEOP structs	No	No	For more information, see VARIABLEOP Struct below.
serializations	List of SERIALIZE structs	No	No	For more information, see SERIALIZE Struct below.
comment	String	No	No	Adds information about the steps.
wapi	String	No	No	The WAPI version. When this is set, the username and password (auth username and auth password) specified for the endpoint are ignored for WAPI related steps only. Other steps still use the configured auth username and auth password. All WAPI requests are sent to the Grid Master IP. The path of the requests is <code>/wapi/[version]</code> with the appended path as specified in the step. The override path option as well as any path configured in the session management template will have no effect.
wapi_quoting	ENUM	No	No	The valid value is one of the following: JSON or XML. If this is not specified, it is set to JSON. The default is JSON. JSON and XML parsing performs as usual for WAPI.
dxl_topic	String	No	Path only	The DXL topic that is used to send data by DXL. <b>Note:</b> This variable is applicable only for the DXL templates.

## VARIABLEOP Struct

Note that all fields are optional, except operation and type. No schema validation is made for variable combinations. Validation of the schema is performed when you add or modify the action template., no specific JSON schema validation is made for variable combinations as it would be too complex, the validation will be done at template conversion time (which is still done when users upload / edit templates).



### Note

All variables listed in the following table are applicable only for the VARIABLEOP operation in the STEP struct. The VARIABLEOP struct converts specified keys/values in a complex variable (such as LIST, DICTIONARY, or XML values with attributes) or performs operations on an existing complex variable. No schema validation is made for variable combinations. Validation of the schema is performed when you add or modify the action template.

## VARIABLEOP Struct

Variable	Type	Mandatory	Sub	Description
source	String	No	Yes	You must specify the variable name, including the namespace. You can include nested variables. Example: L:DICTIONARY{key}{list}
destination	String	No	Yes	You must specify the variable name, including the namespace. You can include nested variables. Example: L:DICTIONARY{key}{list}
operation	ENUM	Yes	Yes	Specifies the operation being performed for the value defined or created in type. The valid value is one of the following: PUSH, POP, SHIFT, UNSHIFT, or ASSIGN. Note that both operation and type are mandatory for the struct. For examples, see Using OPERATION in the VARIABLE Struct below.

Variable	Type	Mandatory	Sub	Description
type	ENUM	Yes	Yes	<p>Specifies the type of variables for the operation. Valid value is one of the following: LIST, DICTIONARY, COMPOSITE, or SINGLE. This field defines how the values in <code>keys</code>, <code>values</code>, and <code>composite_value</code> are combined to create a variable while <code>operation</code> defines how the variable is being used. The appliance handles <code>keys</code>, <code>values</code> and <code>composite_value</code> differently based on the type you select, as follows:</p> <ul style="list-style-type: none"> <li>• LIST: <code>keys</code> will be ignored; <code>values</code> will contain the required values and be serialized in a simple list of strings.</li> <li>• DICTIONARY: The list of <code>key/value</code> pairs in <code>keys</code> and <code>values</code> will be put in a simple dictionary.</li> <li>• COMPOSITE: The list of <code>key/value</code> pairs in <code>keys</code> and <code>values</code> will be put in an XMLA dictionary with a name set to <code>name</code> and a value set to either the value in <code>composite_value</code> OR the value in the variable set in <code>source</code>. See Using the COMPOSITE Type below for details about how to use COMPOSITE.</li> <li>• SINGLE: <code>keys</code> will be ignored; the value will be the first element in the <code>values</code> list.</li> </ul>
name	String	No	Yes	This is applicable only if you select COMPOSITE as the variable type. Enter the name value for this type.
keys	List of strings	No	Yes	This can contain arbitrary serializable constructs.
values	List of strings	No	Yes	This can contain arbitrary serializable constructs. When you specify this field inside any variable operation for the VARIABLEOP step type, it is possible to use the constant of a specific type, instead of using string. For details about how to <code>specific type</code> , instead of using string. For details about how to use constant, see Constant Specification for values below.
composite_value	String	No	Yes	This can contain arbitrary serializable constructs.
destination_key	String	No	Yes	This can contain arbitrary serializable constructs.
if_exists	ENUM	No	Yes	The action taken if the variable exists. The valid value is one of the following: SKIP, ERROR, or NEXT.

Variable	Type	Mandatory	Sub	Description
if_exists_next	String	No	Yes	Enter the name for the next step.
condition	CONDITION Struct	No		This can be specified for POP and UNSHIFT operations.

### Using the COMPOSITE Type

When you use COMPOSITE as the type in the VARIABLE struct, the list of key/value pairs in `keys` and `values` is put in an XMLA dictionary with a name set to `name` and a value set to either the value in `composite_value` OR to the value in the variable set in `source`.

Here is an example of using COMPOSITE as the type when you use the following variables and values:

Field	Value
type	COMPOSITE
name	varname
keys	["a1", "a2"]
values	["value is \${L::SOMEVAR}", "2"]
composite_value	"123"

It results in the following:

```
{
  "<xmla>": true,
  "attrs": {
    "a1": "value is 123",
    "a2": "2",
  },
  "index": {},
  "value": ["123"],
  "name": "varname"
}
```

#### Note

If source is set, the value in source is used as is (this can be seen especially for PUSH operations).

## Using OPERATION in the VARIABLE Struct

### operation ASSIGN

If operation is set to ASSIGN, the value in type will be simply put in the variable identified by `destination`, overwriting any value it might already have.

The following are several examples of the ASSIGN operation having namespace E as follows:

```
{  
  'some_field': 'some_value',  
  'some_list': ['item1', 'item2']  
}
```

#### (1) When assigning a single value:

```
{  
  'operation': 'ASSIGN',  
  'type': 'SINGLE',  
  'destination': 'L:var',  
  'values': ['value']  
}
```

It returns name space L as follows:

```
{'var': 'value'}
```

#### (2) When assigning a list:

```
{  
  'operation': 'ASSIGN',  
  'type': 'LIST',  
  'destination': 'L:list',  
  'values': ['value', '${E:A:some_field}']  
}
```

It returns the following:

```
{'list': ['value', 'some_value']}
```

**(3) When assigning a dictionary:**

```
{  
  'operation': 'ASSIGN',  
  'type': 'DICTIONARY',  
  'destination': 'L:dict',  
  'keys': ['${E:A:some_field}_1', '${E:A:some_field}_2'],  
  'values': ['${E:A:some_list[0]}', '${E:A:some_list[1]}']  
}
```

It returns the following:

```
{'dict': {'some_value_2': 'item2', 'some_value_1': 'item1'}}
```

**(4) When assigning a composite value, as follows:**

```
{  
  'operation': 'ASSIGN',  
  'type': 'COMPOSITE',  
  'destination': 'L:composite',  
  'name': 'tag_name',  
  'keys': ['attr_1', 'attr_2'],  
  'values': ['${E:A:some_field}', 'another_value'],  
  'composite_value': 'tag_content'  
}
```

It gives the following:

```
{'composite': {  
  'index': {}, '<xmla>': True,  
  'attrs': {  
    'attr_2': 'another_value',  
    'attr_1': 'some_value'  
  }  
}
```



```
},  
'value': ['tag_content'],  
'name': 'tag_name'  
}}
```

This composite value can be serialized to the following XML:

```
<tag_name attr_1="some_value" attr_2="another_value">tag_content</tagname>
```

operation PUSH and SHIFT without setting `if_exists`

If `destination_key` is set, the assumption is to push or shift to a dictionary; otherwise, it is to push or shift to a list.

For pushing to a list, if the `destination` variable exists and is a list, the list of values in `values` will be added to the existing list. With PUSH, it means that the values will be added to the right side of the list (at the end) of the list. With SHIFT however, the values will be added to the left side (at the beginning) of the list.

If the `destination` variable is a dictionary, it will be converted to a list that contains only one element with the value appended to it. This means that for XML deserializing and JSON values, it is possible to create a list by pushing it on to an existing scalar value.

For pushing to a dictionary, it means that with `destination_key` set, if the `destination` variable exists and is anything but a dictionary, it is considered as an error. If it is a dictionary however, the value above will be added with its key set to `destination_key` overwriting any existing value.

If the variable denoted by `destination` has a composite value, the COMPOSITE type can be pushed or shifted to it only. If the destination composite value contains text, the text will be replaced with the composite value. In other cases, the PUSH operation using the new value will be added to the right side of the children list. For the SHIFT operation, the new value will be added to the left side of the children list.

The following are several examples of the ASSIGN operation having namespace E:

```
{'some_field': 'some_value'}
```

and namespace L as follows:

```
{  
  'some_list': ['item1', 'item2'],  
  'some_dict': {'key': 'val'},  
  'comp1': {  
    'index': {}, '<xmla>': True,  
    'attrs': {  
      'attr1_b': 'another_value',  
      'attr1_a': 'some_value'  
    }  
  }  
}
```

```

    },
    'value': ['tag_1_content'],
    'name': 'tag_1'
  },
  'comp2': {
    'index': {}, '<xmla>': True,
    'attrs': {},
    'value': ['tag_2_content'],
    'name': 'tag_2'
  }
}

```

**(1) When PUSH or SHIFT to a list:**

```

{
  'operation': 'PUSH',
  'type': 'SINGLE',
  'destination': 'L:some_list',
  'values': ['${E:A:some_field}_right'],
},
{
  'operation': 'SHIFT',
  'type': 'LIST',
  'destination': 'L:some_list',
  'values': ['left_${E:A:some_field}']
},

```

It returns namespace L as the following:

```

{'some_list': [['left_some_value'], 'item1', 'item2', 'some_value_right']}

```

**(2) When PUSH or SHIFT to a dictionary, it returns the following:**

```
{'some_dict': {  
  'push1': 'item',  
  'push2': ['item_1', 'item_2'],  
  'key': 'val'  
}}
```

---

**Note:** There is no difference between PUSH and SHIFT when destination is DICTIONARY.

---

**(3) When PUSH or SHIFT to a composite value:**

```
{  
  'operation': 'PUSH',  
  'type': 'COMPOSITE',  
  'destination': 'L:comp1',  
  'name': 'pushed',  
  'composite_value': ''  
},  
{  
  'operation': 'SHIFT',  
  'type': 'COMPOSITE',  
  'destination': 'L:comp1',  
  'source': 'L:comp2'  
},
```

It returns the following:

```
{  
  'comp1': {  
    'index': {'pushed': 1, 'tag_2': 0},  
    '<xmla>': True,
```

```

      'attrs': {'attr1_a': 'some_value', 'attr1_b': 'another_value'},
      'value': [
        {'index': {}, '<xmla>': True, 'attrs': {}, 'value': ['tag_2_content'],
        'name': 'tag_2'},
        {'index': {}, '<xmla>': True, 'attrs': {'attr': 'val'}, 'value': [],
        'name': 'pushed'}
      ],
      'name': 'tag_1'
    }
  }
}

```

Note that the first operation (PUSH) overwrites existing value and the second operation (SHIFT) shifts the value to the existing list. This composite value can be serialized to the following XML:

```

<tag_1 attr1_a="some_value" attr1_b="another_value">
  <tag_2>tag_2_content</tag_2>
  <pushed attr="val"/>
</tag_1>

```

operation PUSH and SHIFT with `if_exist` set

If `destination_key` is set, the assumption is to push or shift to a dictionary. Otherwise it is to push or shift to a list.

If you are pushing to a list, and the variable already exists, the operation specified there will be executed as follows: SKIP will not modify the original value, NEXT will not modify and jump to the step specified in `if_exists_next`, and ERROR will instead return an error.

If you are pushing to a dictionary, and the variable matches, there is no need to modify the values because they would already be the same. Therefore, SKIP will be a no-operation, NEXT will not modify and jump to the step specified in `if_exists_next`, and ERROR will instead return an error.

The decision for whether a new value exists or not is a simple comparison with the value/attributes of the existing variables in the list or dictionary in `destination_key`. They all have to match for `if_exists` for the operation to be triggered.

If variable denoted by `destination` has a composite value, the semantics is as above, with additional checks. The check is complete before the value is added to the list of children. During the check, the new composite value is compared one by one with the existing values. If the values are the same, the "if\_exists" logic is triggered as above.

The following are a few examples using namespace L as follows:

```

{
  'some_list': ['item1', 'item2'],

```

```

'some_dict': {'key1': 'val1', 'key2': 'val2'},
'list_of_lists': [['a', 'b'], ['a', '2'], ['1', '2']],
'composite': {
  '<xmla>': True,
  'index': {'inner_1': 0, 'inner_2': 1},
  'name': 'outer',
  'attrs': {'outer_attr': 'outer_val'},
  'value': [
    {
      '<xmla>': True,
      'index': {},
      'name': 'inner_1',
      'attrs': {},
      'value': ['inner_1_content']
    },
    {
      '<xmla>': True,
      'index': {},
      'name': 'inner_2',
      'attrs': {'inner_attr': 'inner_val'},
      'value': []
    }
  ]
}
}

```

**(1) When PUSH to a list with the item exists:**

```
{
```

```
'operation': 'PUSH',  
'type': 'SINGLE',  
'destination': 'L:some_list',  
'values': ['item2'],  
'if_exists': 'ERROR'  
}
```

It triggers the `'if_exists'` logic. `'ERROR'` means stopping the template execution with error, `'SKIP'` means 'do nothing for this operation' and `'NEXT'` will jump to the step with name specified in `'if_exists_next'`.

### (2) When PUSH to a dictionary with the item exists:

```
{  
'operation': 'PUSH',  
'type': 'SINGLE',  
'destination': 'L:some_dict',  
'destination_key': 'key2',  
'values': ['val2'],  
'if_exists': 'ERROR'  
}
```

It triggers an error. Note that both `key` and `value` should be the same in order to trigger the `'if_exists'` logic. If `key` is the same but `value` differs, `value` will be overwritten. If `key` is different, the new `key` with `value` will be added.

### (3) When PUSH to a list of list with the item exists:

```
{  
'operation': 'PUSH',  
'type': 'LIST',  
'destination': 'L:list_of_lists',  
'values': ['a', '2'],  
'if_exists': 'ERROR'  
}
```

It triggers an error.

**(4a) When PUSH to a composite value with the item exists, as follows:**

```
{  
  'operation': 'PUSH',  
  'type': 'COMPOSITE',  
  'destination': 'L:composite',  
  'name': 'inner_1',  
  'composite_value': 'inner_1_content',  
  'if_exists': 'ERROR'  
}
```

It triggers an error.

**(4b) When PUSH to a composite value with the item exists, as follows:**

```
{  
  'operation': 'PUSH',  
  'type': 'COMPOSITE',  
  'destination': 'L:composite',  
  'name': 'inner_2',  
  'keys': ['inner_attr'],  
  'values': ['inner_val'],  
  'composite_value': '',  
  'if_exists': 'ERROR'  
}
```

It triggers an error.

operation POP and UNSHIFT

If the `source` variable is anything but a dictionary, a list, or a composite value, an error is returned. If `keys` is set, the `source` value is assumed to be a dictionary. Otherwise, an error is returned.

If the `source` variable is a list, a single value is removed from the list. If the operation is POP, it is removed from the right side (from the end). If the operation is UNSHIFT, it is removed from the left (at the beginning). The removed value is put in the `destination` variable.

If the `source` variable is a composite value, the operation is on the children list of the composite value. Note that it is not possible to POP or UNSHIFT from a composite value, if it has text only.

If the `source` variable is a dictionary, the specified `keys` and/or key/value pairs will be removed from the variable in `source`. If only `keys` are passed, those keys will be unconditionally removed. If `keys` and `values` are passed, the key will be removed from the `source` variable only if its value matches what is passed in `values` here. If only one key or key/value pair is passed, the removed value will be put in `destination`. Otherwise, it will simply be dropped.

For the `values` match, the value in the list or dictionary will be serialized with quoting specified in the template before being compared to the string value passed in `values`. This also means that using comparison on large lists would potentially be a slow operation.

In the POP or UNSHIFT case, `type` is validated to be the same as the type of the popped or unshifted variable (you can validate if they were planning to pop a simple scalar if there is one). If it is not, an error will be raised if debug is turned on.

The following are a few examples using namespace L as follows:

```
{
  'some_list': ['item1', 'item2'],
  'some_dict': {'key1': 'val1', 'key2': 'val2'},
  'list_of_lists': [['a', 'b'], ['1', '2']],
  'list_of_dicts': [{'a': '1'}, {'b': '2'}],
  'composite': {
    '<xmla>': True,
    'index': {'inner_1': 0, 'inner_2': 1},
    'name': 'outer',
    'attrs': {'outer_attr': 'outer_val'},
    'value': [
      {
        '<xmla>': True,
        'index': {},
        'name': 'inner_1',
```



```
    'attrs': {},  
    'value': ['inner_1_content']  
  },  
  {  
    '<xmlla>': True,  
    'index': {},  
    'name': 'inner_2',  
    'attrs': {'inner_attr': 'inner_val'},  
    'value': []  
  }  
]  
}  
}
```

**(1) When POP from a list:**

```
{  
  "operation": "POP",  
  "type": "SINGLE",  
  "source": "L:some_list",  
}
```

It gives the variable as follows:

```
{"some_list": ["item1"]}
```

**(2) When UNSHIFT from a list:**

```
{  
  "operation": "UNSHIFT",  
  "type": "SINGLE",  
  "source": "L:some_list",  
}
```

It returns the following:

```
{"some_list": ["item2"]}
```

**(3) When UNSHIFT from a dict by key:**

```
{  
  "operation": "UNSHIFT",  
  "type": "SINGLE",  
  "source": "L:some_dict",  
  "keys": ["key1"]  
}
```

It returns the following:

```
{"some_dict": {"key2": "val2"}} There is no difference between POP/UNSHIFT  
fro dictionarries
```

**(4) When POP from a dict by multiple keys:**

```
{  
  "operation": "POP",  
  "type": "SINGLE",  
  "source": "L:some_dict",  
  "keys": ["key1", "key2"]  
}
```

It returns the following:

```
{"some_dict": {}}
```

**(5) When POP from a dict when a key is absent:**

```
{  
  "operation": "POP",  
  "type": "SINGLE",  
  "source": "L:some_dict",  
  "keys": ["key1", "absent_key"]  
}
```

It returns an error in the DEBUG mode. For non-DEBUG mode, all existing keys are POPed. Note that items are popped one by one, so key1 item is popped before the error is returned.

```
{"some_dict": {"key2": "val2"}}
```

**(6) When POP from composite value:**

```
{  
  "operation": "POP",  
  "type": "SINGLE",  
  "source": "L:composite",  
}
```

It returns the following:

```
{  
  "composite": {  
    "index": {"inner_1": 0},  
    "<xmla>": True,  
    "name": "outer",  
    "value": [{  
      "index": {},  
      "<xmla>": True,  
      "name": "inner_1",  
      "value": ["inner_1_content"],  
      "attrs": {}  
    }],  
    "attrs": {"outer_attr": "outer_val"}  
  }  
}
```

This can be serialized to the following:

```
<outer outer_attr="outer_val">  
  <inner_1>inner_1_content<inner_1>
```

```
</outer>
```

**(7) When conditional POP/UNSHIFT from a list:**

```
{  
  "operation": "POP",  
  "type": "SINGLE",  
  "source": "L:some_list",  
  "values": ["item1"]  
}
```

It returns the following:

```
{"some_list": ["item2"]}
```

There is no difference between POP and UNSHIFT when 'values' is specified. When source is either **LIST** or **COMPOSITE**, all occurrences of a value are deleted.

**(8) When conditional POP/UNSHIFT from a dictionary:**

```
{  
  "operation": "UNSHIFT",  
  "type": "SINGLE",  
  "source": "L:some_dict",  
  "keys": ["key2"],  
  "values": ["val2"]  
}
```

It returns the following:

```
{"some_dict": {"key1": "val1"}}
```

**(9) When conditional POP/UNSHIFT of multiple values:**

```
{  
  "operation": "UNSHIFT",  
  "type": "SINGLE",  
  "source": "L:some_dict",  
  "keys": ["key2", "key2"],
```

```
"values": ["abc", "val2"]
}
```

It returns the following:

```
{"some_dict": {"key1": "val1"}}
```

**(10) When conditional POP/UNSHIFT with list values:**

```
{
  "operation": "POP",
  "type": "SINGLE",
  "source": "L:list_of_lists",
  "values": ["\"['1', '2']\""]
}
```

It returns the following:

```
{"list_of_lists": [["a", "b"]]}
```

For conditional POP/UNSHIFT, non-string values are serialized with the template's `quoting`. For JSON, `quoting` is added to `value`.

**(11) When conditional POP/UNSHIFT with dictionary value:**

```
{
  "operation": "POP",
  "type": "DICTIONARY",
  "source": "L:list_of_dicts",
  "values": ["\"{'b': '2'}\""]
}
```

It returns the following:

```
{"list_of_dicts": [{"a": "1"}]}
```

**(12a) When conditional POP/UNSHIFT with the following composite values:**

```
{
  "operation": "POP",
  "type": "COMPOSITE",
```

```

"source": "L:composite",
"values": [{"index": {}, '<xmla>': True, 'name': 'inner_2', 'value':
[],
'attrs': {'inner_attr': 'inner_val'}}\"]
}

```

It returns the following:

```

{
  "composite": {
    "index": {"inner_1": 0},
    "<xmla>": True,
    "name": "outer",
    "value": [{
      "index": {},
      "<xmla>": True,
      "name": "inner_1",
      "value": ["inner_1_content"],
      "attrs": {}
    }],
    "attrs": {"outer_attr": "outer_val"}
  }
}

```

**(12b) When conditional POP/UNSHIFT with the following composite values, and with the template quoting set to XMLA:**

The example in (12a) can be specified as follows:

```

{
  "operation": "POP",
  "type": "COMPOSITE",
  "source": "L:composite",
  "values": ["<inner_2 inner_attr=\"inner_val\"/>"]
}

```

```
}
```

Type check after POP/UNSHIFT:

```
{
```

```
  "operation": "POP",
```

```
  "type": "LIST",
```

```
  "source": "L:some_dict",
```

```
  "keys": ["key2"]
```

```
}
```

It returns an error in the DEBUG mode. Since the type of POPped value is not **LIST**, the possible values for type are **'SINGLE'**, **'LIST'**, **'DICTIONARY'**, and **'COMPOSITE'**, where **'SINGLE'** means 'no check'. Note that type check is done after the item is retrieved from `source`. In the **DICTIONARY** case, when several keys are specified, there is no type check. Type check is done after value comparison.

**(13) When putting the item to 'destination':**

```
{
```

```
  "operation": "POP",
```

```
  "type": "COMPOSITE",
```

```
  "source": "L:composite",
```

```
  "destination": "L:sub_item"
```

```
}
```

It returns the following:

```
{
```

```
  "composite": {
```

```
    "index": {"inner_1": 0},
```

```
    "<xmla>": True,
```

```
    "name": "outer",
```

```
    "value": [
```

```
      {"index": {}, "<xmla>": True, "name": "inner_1", "value":
```

```
      ["inner_1_content"],
```

```
      "attrs": {}}
```

```

    ],
    "attrs": {"outer_attr": "outer_val"}
  },
  "sub_item": {
    "index": {},
    "<xmla>": True,
    "name": "inner_2",
    "value": [],
    "attrs": {"inner_attr": "inner_val"}
  }
}

```

If the value popped/unshifted from source and destination is specified, the value is written to the destination. The value is written after the type check.

The L:sub\_item can be serialized as the following:

```
<inner_2 inner_attr="inner_val"/>
```

### Constant Specification for values

When you specify the `values` field inside any variable operation for the VARIABLEOP step type, it is possible to use the constant of a specific type, instead of string.

For example, the following describes three string values:

```
'values': ['True', '42', '${L:some_var}']
```

If you specify `values` as `{X:something}`, `'something'` is represented as a constant of type X. Note that this is the only allowed syntax. For example, `text{X:something}` is treated as a string, instead of a constant specification.

The following types are supported in constant specification:

- **Bool type (B):** Both `{B:1}` and `{B:true}` are evaluated as bool true, while `{B:0}` and `{B:false}` are evaluated as bool false. Note that both 'true' and 'false' are not case-sensitive. Therefore, True, FALSE or even tRuE are all allowed. Other invalid values such as 'B:' and `{B:anythingelse}` are evaluated as bool=false (the 'default value') without further debugging; and an error is logged in the debug mode.
- **Integer type (I):** `{I:}` is evaluated to its corresponding integer value. For example, `{I:-42}` is evaluated as integer -42. The default value is 0.
- **Empty type (E):** `{E:}` is evaluated as None. `{E:[]}` is evaluated as an empty list. `{E:{}}` is evaluated as an empty dictionary. The default value is None.
- **Float type (F):** The float type is similar to the integer type. Float numbers can be specified as `{F:0.0}`, `{F:-4.2}`, `{F:-4E+2}` or `{F:+1.23e-45}`. The default value is 0.0.
- **String type (S):** `{S:}` is evaluated as string ". For example, `{S:{I:42}}` is string '{I:42}'. There is no default value for string. For the unknown type (e.g. `{U:%^&}`), the default value is an empty string.



The following example illustrates how to use constant specification to define `ipv4addr` with a DHCP option:

```
{
  "ipv4addr": "1.2.3.4",
  "options": [
    {
      "name": 'dhcp-lease-time',
      "num": 51,
      "use_option": true,
      "value": '43200',
      'vendor_class': "DHCP"
    }
  ]
}
```

Note that `num` is an integer (in this string: `"num": "51"`) and `use_option` is bool (in this string: `"use_option": "true"`).

In the action template, you can do the following:

```
{
  "steps": [
    {
      "operation": 'VARIABLEOP',
      "variable_ops": [
        {
          "operation": "ASSIGN",
          "type": "DICTIONARY",
          "destination": "L:addr",
          "keys": ["ipv4addr", "options"],
          "values": ["1.2.3.4", "{E:[]}"]
        }
      ],
    }
  ],
}
```

```

{
  "operation": "PUSH",
  "type": "DICTIONARY",
  "destination": "L:addr{options}",
  "keys": [ "name", "num", "use_option", "value", "vendor_class"],
  "values": ["dhcp-lease-time", "{I:51}", "{B:true}", "43200",
"DHCP"]
}
],
{
  "operation": "PUT",
  "transport": {"path": "record:host_ipv4addr"},
  "body": "${L:J:addr}"
}
]
}

```

## SERIALIZE Struct

### *SERIALIZE Struct*

**Note:** All variables listed in the following table are applicable only for the SERIALIZE operation in the STEP struct. The SERIALIZE structs inside the step are executed in sequence.

Field	Type	Mandatory	Sub	Description
content	String	Yes	Yes	Defines what to serialize. This can contain arbitrary variables and text.
destination	String	Yes	Yes	Defines the destination for the serialization.

## TRANSPORT Struct

### *TRANSPORT Structs*

Variable	Type	Mandatory	Sub	Description
path	String	No	Yes, if specified in step.	If present, it is appended to the endpoint URI.
override_path	Boolean	No	No	If this is true, the specified path completely overrides the endpoint URI after the first <i>/</i> .
content_type	String	No	No	If specified, this overrides the endpoint content type.

## RESULT Struct

Note the following:

- **codes** and **regex** are ignored if the operation is SLEEP or NOP.
- At least one of the following is required: **next**, **stop**, or **error**.

### *RESULT Structs*

Variable	Type	Mandatory	Sub	Description
codes	string separated by commas	No	No	The http return code.
regex	REGEX	No	No	If specified, REGEX is matched against the returned body.
next	id	No	No	The step to execute the next template if the code (and REGEX, if specified) match. But if the next template is already executed once, the appliance displays an error.
stop	Boolean	No	No	If set, the execution of the script is stopped. Note that next, error, and stop are mutually exclusive.
error	Boolean	No	No	If set, the execution of the script is stopped and then rerun with an error status if the <code>retry_template</code> is not set to 0. Note that next, error, and stop are mutually exclusive.

## PARAMETER Struct

### *PARAMETER Structs*

Variable	Type	Mandatory	Sub	Description
name	String	Yes	No	Name of the parameter.
value	String	Yes	Yes	Value of the parameter. Note that this value is used as-is, so any strings that are not part of the variables must be URL encoded (%20 for spaces, and etc.)

## CONDITION Struct

### *CONDITION Structs*

Variable	Type	Mandatory	Sub	Description
condition_type	String	Yes	No	This can be one of the following: AND, OR, NAND, or NOR. NAND means not (st1, st2, etc.)
next	String	No	No	The name of the step to jump to if the condition is successful.
eval	String	No	Yes	This is executed if the condition is successful. Generally, it will be an XC: set of operations.
else_eval	String	No	Yes	This is executed if the condition is NOT successful. Generally, it will be an XC: set of operations.
else_next	String	No	Yes	The execution will jump to specified step if the condition is NOT successful.
else_stop	String	No	Yes	The template execution will be stopped if the condition is NOT successful.
else_error	String	No	Yes	Generates an error if the condition is NOT successful.
stop	Boolean	No	No	If the condition is successful the execution will be stopped without any error.
error	Boolean	No	No	If the condition is successful the execution will be stopped with an error.
statements	List of STATEMENT structs	Yes	Yes	The statements to evaluate.

## STATEMENT Struct

### *STATEMENT Structs*

Variable	Type	Mandatory	Sub	Description
left	String	Yes	Yes	The left operand. Note that it is acceptable to include variables that do not exist on the left side (to test if a variable is set in an event). Any error during evaluation of the `left` is ignored even if it is run under debug mode.
op	ENUM	Yes	No	The operation to execute. This can be one of the following: =, >, <, >=, <=, =~, and !~. Note that the >, <, >=, and <= operations try to convert the operands to numbers before executing the comparison, then =~, and !~ are REGEX matches and the right side is considered the REGEX.
right	String	Yes	Yes	The right operand.

## Action Template Variables and Name Spaces

Action templates can access variables in several different name spaces. The following are the available name spaces:

- **C**: http cookies. It supports only the DEL operation (primarily for logout purposes), but it can be used as a substitution origin.
- **Read-Only E**: Event data.
- **H**: http headers. Note that the assigned variables are sent in the next HTTP request and it survives the template execution.
- **Read-Only I**: Template instance variables. It is set in the GUI during the creation of the filter and it also includes the endpoint variables that are set in the GUI when creating an endpoint. Note that the instance variables can override endpoint variables, if needed.
- **L**: Local template variables. This name space is empty at name space startup and will not survive the template invocation.
- **Read-Only P**: Previous endpoint response values (if parsing is enabled for the response.)
- **Read-Only R**: Previous endpoint request http-specific return values. This includes RC, the http status code of the previous request (Example: 200), BODY, and the body of the response.
- **Read-Only RH**: Previous endpoint request that returned http headers.
- **S**: Endpoint session state variables. These variables survive the template invocation (it is used similar to L: name space which is not cleared at the end of the template execution.)
- **Read-Only UT**: Read-only utility variables. The UT: name space contains the following read-only variables:
  - **EPOCH**: EPOCH seconds since Jan 1st, 1970 (integer, 1 second resolution.)
  - **TIME**: UTC time in ISO-8601 format. Example: 2016-04-08T23:09:35Z (1 second resolution).
  - **UUID**: Random UUID of the form 00000000-0000-0000-0000-000000000000.
  - **PROTOCOL**: The protocol used for the request.
  - **URI**: Complete URI used for the request.
  - **HOST**: The endpoint address.
  - **PORT**: The port used for the request.
- **Read-Only XC**: Execute a command on the variable. This results no output.

## Action Template Variable Format

The following sequences are substituted with dynamic data:

```

${<data namespace>:<output format>:<name of the data or operation dependent values>}

```

The sequences cannot contain the following characters as variable names: {, }, [, and ]. Therefore, if names contain keys of dictionary/EA variables, they must be quoted with the \ character. Example: \*\${E::SOME\_EA{name with \{ embedded \} braces or brackets}}\*. But if the \ character is desired, it must be quoted as \. Note that variable substitution happens only once, it is not nested (i.e. if the value contains \${...}, it will not be recursively substituted). Also, the dynamic data is not supported for instance and template variables which are assumed to be immediate values.

The following types of variables allow further qualifiers after the variable name. The qualifiers are always mandatory. Therefore, if a variable is a list or a dictionary, then you must specify [ ] or { }.

- EA (Extensible Attribute) variables or dictionaries. In general, you must specify {EName} for EA variables or dictionaries. For example, if you specify \${E::NETWORK\_EA{City}}, it means the city EA from the network EAs. Note that data specific portion is mandatory for EA variables.
- List variables. You can specify the list index using square brackets. For example, for a list of DNS names if you specify \${E::DNS\_NAMES[0]}, this indicates the first DNS name in the list.

You can also use { } or [ ] to signify the full dictionary or list that is supported only in some encodings. Optionally, you can specify the output format. The following output formats are supported:

- **J**: The output is in the form of JSON formatted variable. It supports deserializing lists as well as dictionaries. Note that strings will have double quotes prepended/appended when serialized with J.
- **j**: The output is in the form of JSON formatted variable. It supports deserializing lists as well as dictionaries without the leading or trailing double quotes, if the variable is a string.
- **X**: The output is in the form of XML formatted variable. It supports only deserializing lists, such as < item >..</item > sequence.
- **U**: The output is in the form of url encoded variable. It supports deserializing lists and the output will be a string (comma separated value.)
- **A**: The output will be a variable, which is as-is.
- **S**: The output is in the form of a string. This is the default for JSON if you do not specify any output format. By default, even numbers will be serialized as JSON strings, meaning the output in a JSON quoted template for a numerical value of 1234 will be "1234".
- **N**: The output is in the form of numbers. For example, if the variable is a boolean, the output will be 0, 1, etc.
- **B**: The output will be a boolean, that is true/false.
- **L**: The length of the variable. This is supported only for lists (the length of the list) and dictionaries (the number of keys).
- **T**: The type of the variable. This can be one of the following characters: 'S' for strings, 'L' for lists, 'D' for dictionaries, 'B' for booleans, 'N' for numbers, and 'O' for otherwise.

The default is set by the template quoting option and by using a variable. If you have not set the template quoting, then JSON will be set, by default. If you have set the template quoting, the output format will be as specified, unless the variable is in the following fields:

- **Headers**: Any variable inside a headers block is serialized by default, as ASIS.
- **Parameters**: Any variable inside a parameters block is serialized by default, as URL.
- **Path**: Any variable inside a path block is serialized by default, as ASIS.

## Command Execution

If you use XC: name space, then the value can be one of the following:

---

**Note:** The argument will be one or two nested variable specifications separated by ':' but without the '\$' symbol. The variables inside 'XC' do not have an output format, so you must use only one ':'.

- 
- **ASSIGN**: Assigns the value to the specified variable. Note that the value assigned is in the format I/S/B:value for integer, string, and boolean values. Example: ASSIGN:variable:value.
  - **DEBUG**: Outputs the specified variable to the debug file (if the log level is not set to DEBUG, this will be ignored), if only the name space is used, the whole name space will be printed.
  - **INC**: Increments the variable value. If the value is not a number, NIOS displays an error.

- **DEC**: Decreases the variable value. If the value is not a number, NIOS displays an error.
- **COPY**: Copies one variable into another. Example: COPY:destination:source.
- **DEL**: Removes the variable. This supports only the C:, H:, L:, and S: name spaces.
- **FORMAT**: Formats the value according to what is specified after the second ':'. Currently, NIOS supports the following formats:
  - **U**: Converts to uppercase value.
  - **L**: Converts to lowercase value.
  - **DATE\_EPOCH**: Assuming that the value is a date expressed in UTC ISO 8601 date format. For example, 2016-03-13T04:50:31Z will be converted to EPOCH seconds.
  - **DATE\_ISO8601**: Assuming that the value contains EPOCH seconds. The value is converted to a date string expressed in UTC ISO 8601 date format. For example, 1467152565 will be converted to 2016-06-28T22:22:45Z. If the variable contains milliseconds, they will be preserved. For example 1467152565.57 will be converted to 2016-06-28T22:22:45.570Z.
  - **DATE\_STRFTIME**: Assuming that the variable contains EPOCH seconds. The value is converted to a date string with the specified format which is passed as the second parameter to the function.
  - **PUNYCODE\_TO\_UTF-8**: Assuming that the variable contains a punycode encoded domain name. The domain name representation will be converted to UTF-8 characters. Note that there might be a failure if the domain name has non-UTF-8 characters in its wire format.
  - **TRUNCATE**: Assuming that the variable is a string and it will be truncated as specified. The format is a number (positive or negative) followed by the letter 'l' or 'r'. The number is the starting character of the string (positive will be counted from the beginning, negative will be counted from the end) and f/t defines if the characters are from, after, or to that point. For example, if a string is 12345, then 1f will produce 2345, 1t will produce 1, -1f will produce 5 and -1t will produce 1234.

Format operations will function like other operations if an error occurs, but the variable is not modified. However, the error can be ignored if the log setting is not set to **Debug**. For information about how to set the logging level, see [Configuring Outbound Endpoints](#).

The following are some examples of using XC operations to increment and decrement IP address strings, create a network range, or remove a specific IP address.

### Incrementing or Decrementing IP addresses

Use the `XC:INC` and `XC:DEC` operations respectively to increment and decrement IP address strings.

Examples:

- For namespace `{'L':{'ip_str': '1.2.3.4'}}`, an evaluating variable `"${XC:DEC:{L:ip_str}}"` results in `{'L':{'ip_str': '1.2.3.3'}}`.
- The same goes for IPv6 addresses. For namespace `{'L':{'ip_str': '2001:db8::2'}}`, an evaluating variable `"${XC:INC:{L:ip_str}}"` results in `{'L':{'ip_str': '2001:db8::3'}}`, and an evaluating variable `"${XC:DEC:{L:ip_str}}"` results in `{'L':{'ip_str': '2001:db8::1'}}`.

The increment of the last address results in the first address. The decrement of the first address results in the last address.

Examples:

INC of '255.255.255.255' results in '0.0.0.0'

INC of 'ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff' results in '::'

DEC of '0.0.0.0' results in '255.255.255.255'

DEC of '::' results in 'fff:ffff:ffff:ffff:ffff:ffff:ffff:ffff'

## Creating Network Ranges from Strings

Use the `XC:NETWORKTORANGE` operation to create a range from a network string. This operation should be specified as the following: `XC:NETWORKTORANGE:{"var_with_network"}:{"var_for_range"}` where network is a string such as `'1.2.3.4/16'` or `'2001:db8:ce4::/48'`

The resulting range is an XMLA value such as the following:

```
{
  '<xmla>': True,
  'name': 'range',
  'attrs': {'from': '1.2.0.0', 'to': '1.2.255.255'},
  'value': [],
  'index': {}
}
```

or

```
{
  '<xmla>': True,
  'name': 'range',
  'attrs': {'from': '2001:db8:ce4::', 'to':
'2001:db8:ce4:ffff:ffff:ffff:ffff:ffff'},
  'value': [],
  'index': {}
}
```

---

**Note:** `"var_with_network"` must be a top-level variable in a namespace such as `E:var`, but not `L:var{key}[42]`.

---

When namespace E is `{'E': {'net': '1.2.3.4/16'}}` and the response is as follows:

```
<SiteConfigResponse success="1">
  <Site id="42" name="this_site" description="testing site" riskfactor="1.0"
isDynamic="0">
    <Description>testing site</Description>
    <Hosts>
      <host>a.com</host>
    </Hosts>
```



```

<Credentials></Credentials>
<Alerting></Alerting>
<ScanConfig configID="21" name="CIS" templateID="cis"
engineID="3" configVersion="3">
  <Schedules></Schedules>
</ScanConfig>
</Site>
</SiteConfigResponse>

```

Then using the following steps:

```

{
  "name": "copy",
  "operation": "NOP",
  "body_list": [
    "${XC:COPY:{L:Site}:{P:PARSE{SiteConfigResponse}}}",
    "${XC:NETWORKTORANGE:{E:net}:{L:range}}"
  ]
},
{
  "name": "add",
  "operation": "VARIABLEOP",
  "variable_ops": [
    {
      "operation": "PUSH",
      "type": "COMPOSITE",
      "destination": "L:Site{Site}{Hosts}",
      "source": "L:range"
    }
  ]
},
{
  "name": "post",
  "operation": "POST",
  "body_list": [
    '<?xml version="1.0" encoding="UTF-8"?>',
    '<SiteSaveRequest session-id="some_sess_id">',
    '${L:x:Site}',

```

```
'</SiteSaveRequest>'
]
}
```

returns the following XML results:

```
<?xml version="1.0" encoding="UTF-8"?>
<SiteSaveRequest session-id="some_sess_id">
  <Site name="this_site" isDynamic="0" description="testing site"
riskfactor="1.0" id="42">
    <Description>testing site</Description>
    <Hosts>
      <host>a.com</host>
      <range from="1.2.0.0" to="1.2.255.255"/>
    </Hosts>
    <Credentials/>
    <Alerting/>
    <ScanConfig configID="21" configVersion="3" name="CIS" templateID="cis"
engineID="3">
      <Schedules/>
    </ScanConfig>
  </Site>
</SiteSaveRequest>
```

### Removing Specific IP Addresses from Hosts or Ranges

Use the `XC:REMOVEIP` operation to remove a specific IP address from a list of hosts or ranges. This operation should be specified as the following: `XC:REMOVEIP:{"var_with_ip"}:{"var_with_list"}` where IP address is a string such as `'1.2.3.4'` or `'2001:db8:ce4::42'`.

The list of hosts/ranges is as follows:

```
[
  {
    '<xmla>': True,
    'name': 'range',
    'attrs': {'from': '1.2.3.1', 'to': '1.2.3.7'},
    'value': [],
    'index': {}
  },
  {
```

```
'<xmlla>': True,  
'name': 'host',  
'attrs': {},  
'value': ['a.com'],  
'index': {}  
}  
]
```

---

**Note:** “var\_with\_IP” must be a top-level variable in a namespace such as `E:var`, but not `L:var{key}` [42].

---

Having namespace E as `{‘E’: {‘ip’: ‘1.2.3.4’}}` and Rapid7 response as follows:

```
<SiteConfigResponse success="1">  
  <Site id="42" name="this_site" description="testing site" riskfactor="1.0" i  
sDynamic="0">  
    <Description>testing site</Description>  
    <Hosts>  
      <host>a.com</host>  
      <range from="1.2.3.1" to="1.2.3.7"/>  
      <host>b.com</host>  
    </Hosts>  
    <Credentials></Credentials>  
    <Alerting></Alerting>  
    <ScanConfig configID="21" name="CIS" templateID="cis"  
engineID="3" configVersion="3">  
      <Schedules></Schedules>  
    </ScanConfig>  
  </Site>  
</SiteConfigResponse>
```

Use the following steps:

```
{  
  "name": "copy",  
  "operation": "NOP",  
  "body_list": [  

```

```

    "${XC:COPY:{L:Site}:{P:PARSE{SiteConfigResponse}}}",
    "${XC:REMOVEIP:{E:ip}:{L:Site{Hosts}}}"
  ]
},
{
  "name": "post",
  "operation": "POST",
  "body_list": [
    '<?xml version="1.0" encoding="UTF-8"?>',
    '<SiteSaveRequest session-id="some_sess_id">',
    '${L:x:Site}',
    '</SiteSaveRequest>'
  ]
}

```

It returns the following XML results:

```

<?xml version="1.0" encoding="UTF-8"?>
<SiteSaveRequest session-id="some_sess_id">
  <Site name="this_site" isDynamic="0" description="testing site"
riskfactor="1.0" id="42">
  <Description>testing site</Description>
  <Hosts>
    <host>a.com</host>
    <range from="1.2.3.1" to="1.2.3.3"/>
    <range from="1.2.3.5" to="1.2.3.7"/>
    <host>b.com</host>
  </Hosts>
  <Credentials/>
  <Alerting/>
  <ScanConfig configID="21" configVersion="3" name="CIS" templateID="cis"
engineID="3">
  <Schedules/>
  </ScanConfig>
</Site>
</SiteSaveRequest>

```

## XC:Remove\* Operations

There are several XC:REMOVE\* operations that can be used to remove one or more IP specified addresses from a list of hosts and ranges. Such list is the deserialized representation of the 'Hosts' element of the Rapid7 response or request.

For example,

```
[
  {
    '<xmla>': True,
    'name': 'host',
    'attrs': {},
    'value': ['a.com'],
    'index': {}
  },
  {
    '<xmla>': True,
    'name': 'range',
    'attrs': {'from': '1.2.3.1', 'to': '1.2.3.7'},
    'value': [],
    'index': {}
  },
  {
    '<xmla>': True,
    'name': 'host',
    'attrs': {},
    'value': ['b.com'],
    'index': {}
  }
]
```

Represents the following `Hosts` element:

```
<Hosts>
<host>a.com</host>
<range from="1.2.3.1" to="1.2.3.7"/>
<host>b.com</host>
</Hosts>
```

Items in the list are iterated one by one upon execution of any REMOVE command. The first argument of the REMOVE\* command is the variable specification that contains an IP address. The second argument is the variable with the list of

ranges or hosts. If the item is 'range', then the specified IP address is deleted from the range. Execution of any REMOVE command may result in no range, one range or two ranges that do not contain any specified IP addresses. The IP addresses specified in non-range items and ranges with different IP addresses are ignored.

- The `XC:REMOVEIP` command specifies the single IP address as string (for example, '1.2.3.4' or '2001:db8:ce4::42').
- The `XC:RMEOVERANGE` command specifies the set of IP addresses as Rapid7 range.
- The `XC:REMOVENET` command specifies the set of IP addresses as network string.

If the first argument denotes an incorrect IP address (range or network respectively), the template execution is stopped and an error is generated for DEBUG mode, and nothing happens for non-DEBUG mode.

#### XC:Parse\* Operations

The XC:PARSE command parses variable data based on the configuration. Some endpoints may return mixed outputs that cannot be parsed using a single parse statement. Hence, Infoblox provides you the flexibility to parse a response or a variable multiple times.

#### Using the XC:PARSE command to evaluate a string

You can use the ``${XC:PARSE:{L:config}:{L:data}}`` command to parse a variable data based on the configuration. Infoblox supports parsing methods such as REGEX, REGEXSPLIT, XML, and so on. The REGEXSPLIT parsing method can be used to split a string to array of items.

The second argument `(L:data)` is updated to the result of parsing operation.

Having the L namespace as:

```
{
"config": {
"parse": "REGEXSPLIT",
"regex": "/"
},
"data": "infoblox.com/contacts"
}
```

The ``${XC:PARSE:{L:config}:{L:data}}`` variable will update `L:data` to `list ["infoblox.com", "contacts"]`.

#### Using the XC:IS\_IP\_IN command to evaluate a string

You can use the `XC:IS_IP_IN` command to evaluate if a string is within a range or a network. The ``${XC:IS_IP_IN:{L:address}:{L:range}}`` variable evaluates the string to `true`. You can specify the second argument as a plain string that contains a mixed list of ranges, or networks or addresses that are separated by a comma. Example: `"10.0.0.1-10.1.0.0,10.1.1.0/24,10.2.2.2"`.

Having the L namespace as:

```
{
"address": "1.2.3.4",
```

```
"range": "1.2.0.0-1.2.255.255"
```

```
}
```

### Using the XC:REMOVEIP command to remove a string

You can remove an IP address from a range as a plain string using the `XC:REMOVEIP` command. Example:

"10.0.0.1-10.0.0.4". In this case, the result format is a plain string with ranges that are separated by a comma. You can specify the second argument as a plain string that contains a mixed list of ranges, or networks, or addresses separated by a comma. Example: "10.0.0.1-10.1.0.0,10.1.1.0/24,10.2.2.2".

Having the L namespace as:

```
{
```

```
"address": "1.2.3.4",
```

```
"range": "1.2.3.1-1.2.3.3,1.2.3.5-1.2.3.20"
```

```
}
```

The ``${XC:REMOVEIP:${L:address}:${L:range}}`` variable updates `L:range` to string "1.2.3.1-1.2.3.3,1.2.3.5-1.2.3.20".

### Using the XC:REMOVENET command to remove a string

You can use the `XC:REMOVENET` command to remove all the IP addresses of a specified network from the list of ranges. Example: "10.0.0.1-10.0.0.4". In this case, the result format is a plain string with ranges that are separated by a comma. You can specify the second argument as a plain string that contains a mixed list of ranges, or networks, or addresses separated by a comma. Example: "10.0.0.1-10.1.0.0,10.1.1.0/24,10.2.2.2".

Having the L namespace as:

```
{
```

```
"network": "1.2.3.0/28",
```

```
"range": "1.2.3.1-1.2.3.20"
```

```
}
```

The ``${XC:REMOVENET:${L:network}:${L:range}}`` variable updates `L:range` to string "1.2.3.15-1.2.3.20".

### Using the XC:REMOVERANGE command to remove a string

You can use the `XC:REMOVERANGE` command to remove all the IP addresses of a specified range from the list of ranges. Example: "10.0.0.1-10.0.0.4". In this case, the result format is a plain string with ranges that are separated by a comma. You can specify the second argument as a plain string that contains a mixed list of ranges, or networks, or addresses separated by a comma. Example: "10.0.0.1-10.1.0.0,10.1.1.0/24,10.2.2.2".

Having the L namespace as:

```
{  
  "range1": "1.2.3.2-1.2.3.10",  
  "range2": "1.2.3.1-1.2.3.20"  
}
```

The `XC:REMOVERANGE:{L:range1}:{L:range2}` variable updates `L:range2` to string `"1.2.3.1,1.2.3.11-1.2.3.20"`.

### Evaluating IP Address in a Range or Network

Use the `XC:IS_IP_IN` operation to validate if an IP address is in the range and network. This operation should be specified as follows: `XC:IS_IP_IN:{var_with_IP}:{var_with_net_or_range}` where IP address is a string such as `'1.2.3.4'` or `'2001:db8:ce4::42'`.

Having the L namespace as:

```
{  
  'ip1': '1.2.3.4',  
  'ip2': '2001:db8:ce5::42',  
  'range': {  
    '<xmla>': True,  
    'name': 'range',  
    'attrs': {'from': '1.2.0.0', 'to': '1.2.255.255'},  
    'value': [],  
    'index': {}  
  },  
  'net': '2001:db8:ce4::/48'  
}
```

The `XC:IS_IP_IN:{L:ip1}:{L:range}` variable will be evaluated to string as `"true"`.

The `XC:IS_IP_IN:{L:ip2}:{L:net}` variable will be evaluated to string as `"false"`.

### Creating Keys

Use the `XC:KEYS` operation to create keys. This operation should be specified as follows: `XC:KEYS:{var_with_dict}:{var_for_key_list}`. It creates a list from the keys of the `var_with_dict` variable and includes it in the `var_for_key_list` variable.

Having L namespace as

```
{
```



```
'some_dict': {  
  'key1': 'value',  
  'key2': 42,  
  'key3': ['item1', 'item2']  
}
```

After evaluating the ``${XC:KEYS:{L:some_dict}:{L:key_list}}`` variable, the L namespace will contain new variable `key_list` with the following values:

```
[  
  'key1',  
  'key2',  
  'key3'  
]
```

## Variable Examples

The following are variable examples.

---

**Note:** The XC: examples refer to the operations discussed in the Command Execution section.

---

### Variables and Results

Variable: ``${E::FQDN}``

Result: Substitute the FQDN value from the event.

Variable: ``${!::QUARANTINE}``

Result: Substitute the quarantine value from the template instance.

Variable: ``${XC:COPY:{L:epoch_timestamp}:{E:timestamp}}`${XC:FORMAT:DATE_EPOCH:{L:epoch_timestamp}}``

Result: Copy the event timestamp value into a local `epoch_timestamp` variable.

Variable: ``${XC:COPY:{L:custom_timestamp}:{E:timestamp}}`${XC:FORMAT:DATE_STRFTIME:{L:custom_timestamp}:{%a, %d %b %Y %H:%M:%S}}``

Result: Copy the event timestamp value to a local `custom_timestamp` variable and format it. In this example, the result might be `Wed, 18 Jun 2015 16:13:11`.

Variable: ``${XC:INC:{S:SERIAL}}`${S::SERIAL}``

Result: Increment the state 'SERIAL' value and substitute its value

Variables and Results
Variable: <code>#{XC:COPY:#{S:SERIAL};{H:X-customheader}}</code> Result: Copy the state SERIAL value to a custom HTTP header for future requests
Variable: <code>#{XC:DEL:#{H:X-loginvalue}}</code> Result: Removes the specified HTTP header.
Variable: <code>#{XC:DEL:#{C:logincookie}}</code> Result: Removes the specified cookie.
Variable: <code>#{XC:ASSIGN:#{L:INT};{I:123}}</code> Result: Assign the integer 123 to the value INT in the L name space.
Variable: <code>#{XC:ASSIGN:#{L:BOOL};{B:true}}</code> Result: Assign true to the value BOOL in the L name space.
Variable: <code>#{XC:ASSIGN:#{L:STR};{S:some } random string}}</code> Result: Assign the string some } random string to the value STR in the L name space.
Variable: <code>#{XC:DEBUG:#{L:INT}}</code> Result: Output the value of the L name space INT variable to the debug file.
Variable: <code>#{XC:DEBUG:#{L:}}</code> Result: Output the values in the whole L name space to the debug file.
Variable: <code>#{XC:FORMAT:TRUNCATE:#{L:VAR};{-1f}}</code> Result: Truncate L:VAR to the last character.

## Event Variables

The E: event name space is populated with variables from the event. For performance reasons, only the variables referred in the template will be available in the name space. This means that if the template is changed and if a new variable is added, it might take some time for the changes to be propagated to all the Grid members and the new variable to be available in future template executions.

You can use `event_type` in an action template to specify the following supported event types: RPZ, LEASE, TUNNEL, NETWORK\_IPV4, NETWORK\_IPV6, RANGE\_IPV4, RANGE\_IPV6, FIXED\_ADDRESS\_IPV4, FIXED\_ADDRESS\_IPV6, HOST\_ADDRESS\_IPV4, HOST\_ADDRESS\_IPV6, and SESSION. Note that SESSION is used only for the login and logout events for the session management templates. For information about action templates, see [Creating Action Templates](#); and for information about session management templates, see [Creating Session Management Templates](#).

The following tables list the supported variables by event type:

### Variables for RPZ Events

#### 45.12 Variables for RPZ Events

NIOS Field Name	Template Variable Name	Supported Filter(s)	Enriched Data	Comment
Timestamp	timestamp (ISO 8601 format)			The timestamp when the event occurred.
Infoblox Member IP	member_ip			Infoblox Member IP (VIP or LAN1) that generated this event.
Infoblox Member Name	member_name			
Action Policy	rpz_policy	equals		(PASSTHRU DROP TCP-ONLY NXDOMAIN NODATA Local-Data CNAME MISS)
RPZ Type	rpz_type	equals		Possible values: BAD, CLIENT-IP, QNAME, IP, NSDNAME, NSIP
Query Name or Query FQDN	query_name	contains, equals, begins with, ends with		
Rule Name	rule_name	contains, equals, begins with, ends with		
Source IP	source_ip	equals, matches range, matches CIDR	Parent range, Associated objects' DNS Names and EAs, Parent Network, Parent Network EA, Discovered data	
Source Port	source_port			
Destination IP	destination_ip			The name server that responded to the RPZ rule.
Query Type	query_type			DNS query type: A, AAAA, CNAME, DNAME, TXT, and other.
Query View Name	query_view_name			Query DNS view name.

NIOS Field Name	Template Variable Name	Supported Filter(s)	Enriched Data	Comment
Event Type	event_type	equals	RPZ LEASE TUNNEL NETWORK_IPV4 NETWORK_IPV6 RANGE_IPV4 RANGE_IPV6 FIXED_ADDRESS_IPV4 FIXED_ADDRESS_IPV6 HOST_ADDRESS_IPV4 HOST_ADDRESS_IPV6 SESSION	SESSION is used for the login and logout events for the session management templates.
<b>Enriched from IPv4 or IPv6 Addresses</b>				
Source IP User Name	ip.username		From IPv4Address and IPv6Address	This field can be empty.
Source IP Associated EA	ip.extattrs		From IPv4Address and IPv6Address pyabs class	Can be empty, EA needs to be stored as "name-value" pair.
Source IP Associated DNS Names	ip.names		From IPv4Address or IPv6Address	List of FQDN
<b>Enriched from Discovered Data</b>				
Source IP Attached Device Model	ip.discovered_data.device_model		From discovered data	
Source IP Attached Device Name	ip.discovered_data.device_port_name		From discovered data	
Source IP Attached Device Port	ip.discovered_data.device_port_type		From discovered data	
Source IP Attached Device Type	ip.discovered_data.device_type		From discovered data	
Source IP Attached Device Vendor	ip.discovered_data.device_vendor		From discovered data	
Source IP Discovered Name	ip.discovered_data.discovered_name		From discovered data	

NIOS Field Name	Template Variable Name	Supported Filter(s)	Enriched Data	Comment
Source IP First Discovered	ip.discovered_data.first_discovered (ISO 8601 format)		From discovered data	
Source IP Discovered MAC	ip.discovered_data.mac_address		From discovered data	
Source IP NetBIOS Name	ip.discovered_data.netbios_name		From discovered data	
Source IP Port Link	ip.discovered_data.port_link_status		From discovered data	
Source IP Port Speed	ip.discovered_data.port_speed		From discovered data	
Source IP Port Status	ip.discovered_data.port_status		From discovered data	
Source IP VLAN Description	ip.discovered_data.port_vlan_description		From discovered data	
Source IP VLAN Name	ip.discovered_data.port_vlan_name		From discovered data	
<b>Enriched from Parent Range</b>				
Source IP Range Start Address	range.start_addr		From parent range	
Source IP Range End Address	range.end_addr		From parent range	
<b>Enriched from Parent Network</b>				
Source IP Network View Name	network.network_view		From parent network	The network view name in string format.
Source IP Network	network.network		From parent network	ip_addr/cidr Example: 1.2.3.4/24
Source IP Network Address	network.ipv4addr		From parent network	ip_addr
Source IP Network Cidr	network.netmask		From parent network	cidr

NIOS Field Name	Template Variable Name	Supported Filter(s)	Enriched Data	Comment
Source IP Network EA	network.extattrs		From parent network	EA name can be any UTF8 characters.
<b>Enriched from Lease Data</b>				
Source IP Lease Start Time	lease.starts (ISO 8601 format)		From lease data	
Source IP Lease End Time	lease.ends (ISO 8601 format)		From lease data	
Source IP Lease State	lease.binding_state		From lease data	Possible values: UNKNOWN, ABANDONED, ACTIVE, BACKUP, DECLINED, EXPIRED, FREE, OFFERED, RELEASED, RESET, STATIC
Source IP Lease Client Host Name	lease.client_hostname		From lease data	
Source IP Lease MAC Address	lease.hardware		From lease data	
Source IP Lease DUID	lease.ipv6_duid		From lease data	
Source IP Fingerprint	lease.fingerprint		From lease data	

### Variables for DHCP Lease Events

When searching for DHCP lease events with associated discovered data, both the "address" and "hardware" or "duid" must match the discovered data. If there is no hardware or DUID, the lease event cannot be associated with any discovered data.

For leases, same IP addresses may be used by multiple systems, so the IP address must match the MAC address or DUID to ensure that the discovered data has the most likely correct value.

### Variables for DHCP Lease Events

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Timestamp	timestamp (ISO 8601 format)			The timestamp when the event occurs.
Infoblox Member IP	member_ip			Infoblox member IP (VIP or LAN1) that has generated the event
Infoblox Member Name	member_name			

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Lease Binding State	binding_state	equals		This can be either UNKNOWN, ABANDONED, ACTIVE, BACKUP, DECLINED, EXPIRED, FREE, OFFERED, RELEASED, RESET, or STATIC.
Lease IP address	address	equals, matches range, matches CIDR	Parent Network EA, Parent Range, DNS Names and EA of associated objects	Some information from parent network is already in the DHCP lease data.
Lease Protocol	protocol			Possible values are IPv4 or IPv6.
Lease Start Time	starts			
Lease End Time	ends			
Client MAC address	hardware			
Client IPv6 DUID	ipv6_duid			
Client Host Name	client_hostname			
Fingerprint		contains, equals, begins with, ends with		
Lease Network View Name	network_view	contains, equals, begins with, ends with		
Lease Network	network			ip_addr/cidr Example: 1.2.3.4/24
Lease Network Address	network_ipaddr			ip_addr
Lease Network CIDR	network_netmask			CIDR
Lease IP Range Start Address	range_start_addr			
Lease IP Range End Address	range_end_addr			

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Event Type	event_type	equals	RPZ LEASE TUNNEL NETWORK_IPV4 NETWORK_IPV6 RANGE_IPV4 RANGE_IPV6 FIXED_ADDRESS_IPV4 FIXED_ADDRESS_IPV6 HOST_ADDRESS_IPV4 HOST_ADDRESS_IPV6 SESSION	SESSION is used for the login and logout events for the session management templates.
<b>Enriched from Network EA</b>				
Lease Network EA	network.extattrs		From parent network EA	EA name can be any UTF-8 characters.
<b>Enriched from IPv4 Address and IPv6 Address</b>				
Lease IP Username	ip.username		From IPv4 Address and IPv6 Address	
Lease IP Associated EA	p.extattrs		From IPv4 Address and IPv6 Address pyabs class	EA must be stored as "<EA name>-<EA value>" pair.
Lease IP Associated DNS Names	ip.names		From IPv4 Address and IPv6 Address	List of FQDN
<b>Enriched from Discovery Data</b>				
Lease IP Attached Device Model	ip.discovered_data.device_model		From discovered data	
Lease IP Attached Device Name	ip.discovered_data.device_port_name		From discovered data	
Lease IP Attached Device Port	ip.discovered_data.device_port_type		From discovered data	
Lease IP Attached Device Type	ip.discovered_data.device_type		From discovered data	
Lease IP Attached Device Vendor	ip.discovered_data.device_vendor		From discovered data	



NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Lease IP Discovered Name	ip.discovered_data.discovered_name		From discovered data	
Lease IP First Discovered	ip.discovered_data.first_discovered (ISO 8601 format)		From discovered data	
Lease IP Discovered MAC	ip.discovered_data.mac_address		From discovered data	
Lease IP NetBIOS Name	ip.discovered_data.netbios_name		From discovered data	
Lease IP Port Link	ip.discovered_data.port_link_status		From discovered data	
Lease IP Port Speed	ip.discovered_data.port_speed		From discovered data	
Lease IP Port Status	ip.discovered_data.port_status		From discovered data	
Lease IP VLAN Description	ip.discovered_data.port_vlan_description		From discovered data	
Lease IP VLAN Name	ip.discovered_data.port_vlan_name		From discovered data	

#### Variables for Object Change Discovery Data

**Note:** Infoblox supports insert and update operations from Network Insight and Cloud Discovery for object change discovery data in the NIOS 8.3.0 release version.

#### Variables for Object Change Discovery Data

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
timestamp	timestamp			Timestamp when event occurs, will be used record_timestamp from incoming data
member_ip	member_ip			Infoblox Member IP

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
member_name	member_name			Infoblox Member Name
object_type	object_type			Object Type, ENUM {DiscoveryData}
operation_type	operation_type	equals {INSERT, MODIFY, DELETE}		Operation Type, ENUM {INSERT, MODIFY, DELETE}
<b>Values</b>				
network	network		Yes	Network (from network ref)
is_ipv4	is_ipv4	equals {true, false}		This boolean flag will indicate whether structure is of type IPv4 or IPv6
ip_address	ip_address	equals, matches range, matches CIDR	Yes (default)	Discovered IP address
mac_address	mac_address		Yes	Discovered MAC address
duid	duid		Yes	DUID associated with IPv6 address
ap_ip_address	ap_ip_address			Discovered IP address of Wireless Access Point
ap_name	ap_name			Discovered name of Wireless Access Point
ap_ssid	ap_ssid			Service set identifier (SSID) associated with Wireless Access Point
bridge_domain	bridge_domain			Discovered bridge domain
cisco_ise_endpoint_profile	cisco_ise_endpoint_profile			Endpoint profile in Cisco ISE
cisco_ise_security_group	cisco_ise_security_group			Name of security group created in Cisco ISE
cisco_ise_session_state	cisco_ise_session_state			session state
cisco_ise_ssid	cisco_ise_ssid			service set identifier
cmp_type	cmp_type			if the IP is coming from a Cloud environment, the Cloud Management Platform type
device_contact	device_contact			Contact information from device on which the IP address was discovered

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
device_location	device_location			Location of device on which the IP address was discovered
device_model	device_model			Model name of the device in the vendor terminology
device_port_name	device_port_name			System name of the interface the IP associates with
device_port_type	device_port_type			Hardware type of the interface the IP associates with
device_type	device_type			Type of the device in the vendor terminology
device_vendor	device_vendor			Vendor name of the device
discovered_name	discovered_name			Name of the ip as seen by the discovery station.
discoverer	discoverer	contains, equals, begins with, ends with	Yes (default)	Name of the discoverer
endpoint_groups	endpoint_groups			List of discovered endpoint groups represented as a single string containing comma-separated values
first_discovered_timestamp	first_discovered_timestamp			When was this ip first seen by the discovery station
iprg_id	iprg_id			iprg(port redundant group) ID of this device interface
iprg_no	iprg_no			Port redundant group no of this device interface
iprg_state	iprg_state			State of this IP address in the group
iprg_type	iprg_type			Type of this prg
is_end_host	is_end_host			Is this object an end host or an infrastructure device for the purpose of discovery
last_discovered_timestamp	last_discovered_timestamp			When was this data discovered
last_updated_timestamp	last_updated_timestamp			Used by NetMRI sync to store timestamp when was this data updated
method	method			What method was being used for network discovery
mgmt_ip_address	mgmt_ip_address			Management IP address of the device if the device has more than one IP

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
netbios_name	netbios_name			Discovered NetBIOS name
network_component_contact	network_component_contact			Contact information from network component on which the IP address was discovered
network_component_description	network_component_description			A descriptive string for the network component
network_component_ip	network_component_ip			IP Address the network component
network_component_location	network_component_location			Location of network component on which the IP address was discovered
network_component_model	network_component_model			Model name of the network component the device is connected to in the vendor terminology
network_component_name	network_component_name			The name of network component
network_component_port_description	network_component_port_description			Description of the port on the network component
network_component_port_id	network_component_port_id			Interface ID of the connected switch/switch-router
network_component_port_name	network_component_port_name			Port name on the network component on which the ip was discovered
network_component_port_number	network_component_port_number			Port number on the network component on which the ip was discovered
network_component_type	network_component_type			The type of network component. Eg. Switch, Router etc.
network_component_vendor	network_component_vendor			Vendor name of the network component the device is connected to
open_ports	open_ports			List of opened ports on the IP address, represented as: "TCP: 21,22,23 UDP: 137,139". Limited to max total 1000 ports
os	os			Guess for OS by network discovery
port_duplex	port_duplex			Duplex settings on the port on the network component
port_link_status	port_link_status			Link Status of the port on the network component
port_speed	port_speed			Speed settings on the port on the network component

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
port_status	port_status			Status of the port on the network component
port_type	port_type			Type of the interface on the network component the device is connected to
port_vlan_description	port_vlan_description			Description of the vlan on the port on the network component
port_vlan_name	port_vlan_name			Name of the vlan on the port on the network component
port_vlan_number	port_vlan_number			Number of the vlan on the port on the network component
task_name	task_name			Name of the task discovers this data
tenant	tenant			Discovered tenant
unmanaged	unmanaged	equals {true, false}		Flag to indicate if the discovered data is associated with a managed object or not
v_adapter	v_adapter			Vmware physical adapter in which entity was found
v_cluster	v_cluster			Vmware cluster in which entity was found
v_datacenter	v_datacenter			Vmware datacenter in which entity was found
v_entity_name	v_entity_name			Vmware entity name in which entity was found
v_entity_type	v_entity_type			Type of encryption to use
v_host	v_host			Vmware host system in which entity was found
v_os	v_os			Vmware OS system in which entity was found
v_switch	v_switch			Vmware virtual switch in which entity was found
vlan_port_group	vlan_port_group			Port group which the virtual machine belongs to
vmhost_ip_address	vmhost_ip_address			IP address of the physical node on which the virtual machine is hosted
vmhost_mac_address	vmhost_mac_address			MAC address of the physical node on which the virtual machine is hosted

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
vmhost_name	vmhost_name			Name of the physical node on which the virtual machine is hosted
vmhost_nic_names	vmhost_nic_names			List of all physical port names used by the virtual switch on the physical node on which the virtual machine is hosted. Represented as: "eth1,eth2,eth3"
vmhost_subnet_cidr	vmhost_subnet_cidr			CIDR subnet of the physical node on which the virtual machine is hosted
vmi_id	vmi_id			ID of the virtual machine
vmi_ip_type	vmi_ip_type			Discovered IP address type
vmi_is_public_address	vmi_is_public_address			Indicates whether the IP address is a public address
vmi_name	vmi_name			Name of the virtual machine
vmi_private_address	vmi_private_address			Private IP address of the virtual machine
vmi_tenant_id	vmi_tenant_id			ID of the tenant which virtual machine belongs to
vport_conf_mode	vport_conf_mode			Configured mode of the network adapter on the virtual switch where the virtual machine connected to
vport_conf_speed	vport_conf_speed			Configured speed of the network adapter on the virtual switch where the virtual machine connected to. Unit is kb
vport_link_status	vport_link_status			Link status of the network adapter on the virtual switch where the virtual machine connected to
vport_mac_address	vport_mac_address			MAC address of the network adapter on the virtual switch where the virtual machine connected to
vport_mode	vport_mode			Actual mode of the network adapter on the virtual switch where the virtual machine connected to
vport_name	vport_name			Name of the network adapter on the virtual switch connected with the virtual machine
vport_speed	vport_speed			Actual speed of the network adapter on the virtual switch where the virtual machine connected to. Unit is kb

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
vswitch_available_ports_count	vswitch_available_ports_count			Numer of available ports reported by the virtual switch on which the virtual machine/vport connected to
vswitch_id	vswitch_id			ID of the virtual switch
vswitch_ipv6_enabled	vswitch_ipv6_enabled			Indicates the virtual switch has IPV6 enabled
vswitch_name	vswitch_name			Name of the virtual switch
vswitch_segment_id	vswitch_segment_id			ID of the network segment on which the current virtual machine/vport connected to
vswitch_segment_name	vswitch_segment_name			Name of the network segment on which the current virtual machine/vport connected to
vswitch_segment_port_group	vswitch_segment_port_group			Port group of the network segment on which the current virtual machine/vport connected to
vswitch_segment_type	vswitch_segment_type			Type of the network segment on which the current virtual machine/vport connected to
vswitch_tep_dhcp_server	vswitch_tep_dhcp_server			DHCP server of the virtual tunnel endpoint (VTEP) in the virtual switch
vswitch_tep_ip	vswitch_tep_ip			IP address of the virtual tunnel endpoint (VTEP) in the virtual switch
vswitch_tep_multicast	vswitch_tep_multicast			Muticast address of the virtual tunnel endpoint (VTEP) in the virtual switch
vswitch_tep_port_group	vswitch_tep_port_group			Port group of the virtual tunnel endpoint (VTEP) in the virtual switch
vswitch_tep_type	vswitch_tep_type			Type of virtual tunnel endpoint (VTEP) in the virtual switch
vswitch_tep_vlan	vswitch_tep_vlan			VLAN of the virtual tunnel endpoint (VTEP) in the virtual switch
vswitch_type	vswitch_type			Type of the virtual switch: standard or distributed

## Variables for Security ADP Events

### *Variables for Security ADP Events*

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment	
Timestamp	timestamp	timestamp (iso 8601 format)			Timestamp when event occurs
Infoblox Member IP	member_ip	member_ip	equals, matches range, matches CIDR		Infoblox Member IP (VIP or LAN1) that has this event generated.
Infoblox Member Name	member_name	member_name	equals		Infoblox Member Name
Rule ID	rule_sid	rule_sid	contains, equals, begins with, ends with		ATP rule ID. String
Rule Name	rule_name	rule_name			ATP rule name
Rule Category	rule_category	rule_category	contains, equals, begins with, ends with		Rule category. String
Rule Log Severity	rule_severity	rule_severity	equals, more severe, less severe		Rule log severity (CRITICAL, MAJOR, WARNING, INFORMATIONAL)
Rule Action	rule_action	rule_action	equal		Rule action (ALERT, DROP, PASS)
Source IP	source_ip	source_ip	equals, matches range, matches CIDR	Parent Range, Associated objects: DNS Names and EA Parent Network, Parent Network EA, Discovery	Source IP
Source Port	source_port	source_port			Source port
Is NAT Client	is_nat_client	is_nat_client			Flag indicating if client is 'nated' (True, False)
NAT First Port	nat_first_port	nat_first_port			Port block start
NAT Last Port	nat_last_port	nat_last_port			Port block end
Query FQDN	query_fqdn	query_fqdn	contains, equals, begins with, ends with		DNS Query FQDN limited by second-level domain name



NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment	
Count of hits	hits_count	hits_count	greater than		Count of ADP hits per interval (15 sec)
<b>Enriched from IPv4 Address or IPv6 Address</b>	<b>WAPI: IPv4 Address / IPv6 Address</b>				
Source IP User Name	ip_username	ip.username		From IPv4Address and IPv6Address	maybe empty
Source IP Associated EA	ip_extattrs	ip.extattrs		From IPv4Address and IPv6Address pyabs class	maybe empty, EA need to be stored as "name-value" pair
Source IP Associated DNS Names	ip_names	ip.names		From IPv4Address or IPv6Address	list of FQDN
<b>Enriched from Discovery data</b>	<b>WAPI: discovery_data struct, inside IPv4 Address / IPv6 Address</b>				
Source IP Attached Device Model	ip_discovered_data_device_model	ip.discovered_data.device_model		From Discovery Data	
Source IP Attached Device Name	ip_discovered_data_device_port_name	ip.discovered_data.device_port_name		From Discovery Data	
Source IP Attached Device Port	ip_discovered_data_device_port_type	ip.discovered_data.device_port_type		From Discovery Data	
Source IP Attached Device Type	ip_discovered_data_device_type	ip.discovered_data.device_type		From Discovery Data	
Source IP Attached Device Vendor	ip_discovered_data_device_vendor	ip.discovered_data.device_vendor		From Discovery Data	
Source IP Discovered Name	ip_discovered_data_discovered_name	ip.discovered_data.discovered_name		From Discovery Data	
Source IP First Discovered	ip_discovered_data_first_discovered	ip.discovered_data.first_discovered (iso 8601 format)		From Discovery Data	
Source IP Discovered MAC	ip_discovered_data_mac_address	ip.discovered_data.mac_address		From Discovery Data	
Source IP NetBIOS Name	ip_discovered_data_netbios_name	ip.discovered_data.netbios_name		From Discovery Data	

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment	
Source IP Port Link	ip_discovered_data_port_link_status	ip.discovered_data.port_link_status		From Discovery Data	
Source IP Port Speed	ip_discovered_data_port_speed	ip.discovered_data.port_speed		From Discovery Data	
Source IP Port Status	ip_discovered_data_port_status	ip.discovered_data.port_status		From Discovery Data	
Source IP VLAN Description	ip_discovered_data_port_vlan_description	ip.discovered_data.port_vlan_description		From Discovery Data	
Source IP VLAN Name	ip_discovered_data_port_vlan_name	ip.discovered_data.port_vlan_name		From Discovery Data	
<b>Enriched from Parent Range</b>	<b>WAPI: range</b>				
Source IP Range Start Address	range_start_addr	range.start_addr		From Parent Range	
Source IP Range End Address	range_end_addr	range.end_addr		From Parent Range	
<b>Enriched from Parent Network</b>	<b>WAPI: network</b>				
Source IP Network View Name	network_network_view	network.network_view		From Parent Network	network view name in string
Source IP Network	network_network	network.network		From Parent Network	ip_addr/cidr, 1.2.3.4/24
Source IP Network Address	network_ip4addr	network.ipv4addr		From Parent Network	ip_addr
Source IP Network Cidr	network_netmask	network.netmask		From Parent Network	cidr
Source IP Network EA	network_extattrs	network.extattrs		From Parent Network EA	EA name can be any utf8 characters.
<b>Enriched from lease data</b>	<b>WAPI: lease</b>				
Source IP Lease Start Time	lease_starts	lease.starts(iso 8601 format)		From Lease Data	

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment	
Source IP Lease End Time	lease_ends	lease.ends(iso 8601 format)		From Lease Data	
Source IP Lease State	lease_binding_state	lease.binding_state		From Lease Data	UNKNOWN, ABANDONED, ACTIVE, BACKUP, DECLINED, EXPIRED, FREE, OFFERED, RELEASED, RESET, STATIC
Source IP Lease Client Host Name	lease_client_hostname	lease.client_hostname		From Lease Data	
Source IP Lease MAC Address	lease_hardware	lease.hardware		From Lease Data	
Source IP Lease DUID	lease_ipv6_duid	lease.ipv6_duid		From Lease Data	
Source IP Fingerprint	lease_fingerprint	lease.fingerprint		From Lease Data	

## Variables for Analytics DNS Tunneling Events

### *Variables for Analytics DNS Tunneling Events*

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Timestamp	timestamp (ISO 8601 format)			When the event occurs
Infoblox Member IP	member_ip			Infoblox member IP (VIP or LAN1) that has generated the event.
Infoblox Member Name	member_name			
Source IP	source_ip	equals, matches range, matches CIDR	Parent Network, Parent Network EA, Discovery	
Domain Name	domain_name			Domain name that was determined as DNS tunneling domain.
RPZ Rule Policy	rpz_policy			RPZ rule policy that was created.

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Event Type	event_type	equals	RPZ LEASE TUNNEL NETWORK_IPV4 NETWORK_IPV6 RANGE_IPV4 RANGE_IPV6 FIXED_ADDRESS_IPV4 FIXED_ADDRESS_IPV6 HOST_ADDRESS_IPV4 HOST_ADDRESS_IPV6 SESSION	SESSION is used for the login and logout events for the session management templates.
Comment	comment			Comment generated from the event by the Analytics system.
<b>Enriched from IPv4 Address or IPv6 Address</b>				
Source IP Username	ip.username		From IPv4 Address and IPv6 Address	
Source IP Associated EA	ip.extattrs		From IPv4 Address and IPv6 Address pyabs class	EA must be stored as "<EA name>-<EA value>" pair.
Source IP Associated DNS Names	ip.names		From IPv4 Address and IPv6 Address	List of FQDN.
<b>Enriched from Discovery Data</b>				
Source IP Attached Device Model	ip.discovered_data.device_model		From discovered data	
Source IP Attached Device Name	ip.discovered_data.device_port_name		From discovered data	
Source IP Attached Device Port	ip.discovered_data.device_port_type		From discovered data	
Source IP Attached Device Type	ip.discovered_data.device_type		From discovered data	
Source IP Attached Device Vendor	ip.discovered_data.device_vendor		From discovered data	
Source IP Discovered Name	ip.discovered_data.discovered_name		From discovered data	

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Source IP First Discovered	ip.discovered_data.first_discovered (ISO 8601 format)		From discovered data	
Source IP Discovered MAC	ip.discovered_data.mac_address		From discovered data	
Source IP NetBIOS Name	ip.discovered_data.netbios_name		From discovered data	
Source IP Port Link	ip.discovered_data.port_link_status		From discovered data	
Source IP Port Speed	ip.discovered_data.port_speed		From discovered data	
Source IP Port Status	ip.discovered_data.port_status		From discovered data	
Source IP VLAN Description	ip.discovered_data.port_vlan_description		From discovered data	
Source IP VLAN Name	ip.discovered_data.port_vlan_name		From discovered data	
<b>Enriched from Parent Range</b>				
Source IP Range Start Address	range.start_addr		From Parent Range	
Source IP Range End Address	range.end_addr		From Parent Range	
<b>Enriched from Parent Network</b>				
Source IP Network View Name	network.network_view		From Parent Network	Network view name in string format.
Source IP Network	network.network		From Parent Network	ip_addr/cidr Example: 1.2.3.4/24
Source IP Network Address	network.ipv4addr		From Parent Network	ip_addr
Source IP Network Cidr	network.netmask		From Parent Network	CIDR
Source IP Network EA	network.extattrs		From Parent Network EA	EA name can be any UTF-8 characters.

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
<b>Enriched from Lease Data</b>				
Source IP Lease Start Time	lease.starts (ISO 8601 format)		From Lease Data	
Source IP Lease End Time	lease.ends (ISO 8601 format)		From Lease Data	
Source IP Lease State	lease.binding_state		From Lease Data	
Source IP Lease Client Host Name	lease.client_hostname		From Lease Data	
Source IP Lease MAC Address	lease.hardware		From Lease Data	
Source IP Lease DUID	lease.ipv6_duid		From Lease Data	
Source IP Fingerprint	lease.fingerprint		From Lease Data	

#### Variables for DB Object Change Event - DHCP Network IPv4

##### *Variables for DB Object Change Event - DHCP Network/Network Container IPv4*

NIOS Field Name	Template Variable Name	Filter	Enriched Data	Comment
Timestamp	timestamp (ISO 8601 format)			When the event occurs
Infoblox Member IP	member_ip			Infoblox member IP (VIP or LAN1) that has generated the event.
Infoblox Member Name	member_name			
WAPI object reference	_ref			
Comment	comment			
Disable	disable	equals		Boolean
Extensible Attributes	extattrs			dictionary of extensible attributes
Network	network	equals, contained in		ip_addr/cidr

NIOS Field Name	Template Variable Name	Filter	Enriched Data	Comment
network_view	network_view	contains, equals, begins with, ends with		String format
Members	members			
MS AD User Data	ms_ad_user_data			
Comment	comment			
network_container	network_container			
options	options			
unmanaged	unmanaged			
Event Type	event_type	equals	RPZ LEASE TUNNEL NETWORK_IPV4 NETWORK_IPV6 RANGE_IPV4 RANGE_IPV6 FIXED_ADDRESS_IPV4 FIXED_ADDRESS_IPV6 HOST_ADDRESS_IPV4 HOST_ADDRESS_IPV6 SESSION	SESSION is used for the login and logout events for the session management templates.

#### Variables for DB Object Change Event - DHCP Network IPv6

#### *Variables for DB Object Change Event - DHCP Network/Network Container IPv6*

NIOS Field Name	Template Variable Name	Filter	Enriched Data	Comment
Timestamp	timestamp (ISO 8601 format)			When the event occurs
Infoblox Member IP	member_ip			Infoblox member IP (VIP or LAN1) that has generated the event.
Infoblox Member Name	member_name			
WAPI object reference	_ref			
Comment	comment			

NIOS Field Name	Template Variable Name	Filter	Enriched Data	Comment
Disable	disable	equals		Boolean
Extensible Attributes	extattrs			Dictionary of extensible attributes
Network	network	equals, contained in		ip_addr/cidr
network_view	network_view	contains, equals, begins with, ends with		String
Members	members			
MS AD User Data	ms_ad_user_data			
Comment	comment			
network_container	network_container			
options	options			
unmanaged	unmanaged			
Event Type	event_type	equals	RPZ LEASE TUNNEL NETWORK_IPV4 NETWORK_IPV6 RANGE_IPV4 RANGE_IPV6 FIXED_ADDRESS_IPV4 FIXED_ADDRESS_IPV6 HOST_ADDRESS_IPV4 HOST_ADDRESS_IPV6 SESSION	SESSION is used for the login and logout events  for the session management templates.

#### Variables for DB Object Change Event - DHCP Range IPv4

#### Variables for DB Object Change Event - DHCP Range IPv4

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Timestamp	timestamp (ISO 8601 format)			When the event occurs
Infoblox Member IP	member_ip			Infoblox Member IP (VIP or LAN1) that has generated the event.



NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Infoblox Member Name	member_name			
WAPI object reference	_ref			
Comment	comment			
Disable	disable	equals		Boolean
Extensible Attributes	extattrs			Dictionary of extensible attributes
Network	network	equals, contained in		ip_addr/cidr
network_view	network_view	contains, equals, begins with, ends with		String
Member	member			
Boot File	bootfile			
Start Address	start_addr			
End Address	end_addr			
MAC Filter Rules	mac_filter_rules			
MS AD User Data	ms_ad_user_data			
Next Server	nextserver			
Server Association Type	server_association_type	contains, equals, begins with, ends with		

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Event Type	event_type	equals	RPZ LEASE TUNNEL NETWORK_IPV4 NETWORK_IPV6 RANGE_IPV4 RANGE_IPV6 FIXED_ADDRESS_IPV4 FIXED_ADDRESS_IPV6 HOST_ADDRESS_IPV4 HOST_ADDRESS_IPV6 SESSION	SESSION is used for the login and logout events for the session management templates.

#### Variables for DB Object Change Event - DHCP Range IPv6

##### *Variables for DB Object Change Event - DHCP Range IPv6*

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Timestamp	timestamp (ISO 8601 format)			Timestamp when the event occurs.
Infoblox Member IP	member_ip			Infoblox Member IP (VIP or LAN1) that has generated the event.
Infoblox Member Name	member_name			
WAPI object reference	_ref			
Comment	comment			
Disable	disable	equals		Boolean
Extensible Attributes	extattrs			Dictionary of extensible attributes
Network	network	equals, contained in		ip_addr/cidr
network_view	network_view	contains, equals, begins with, ends with		String
Member	member			

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Boot File	bootfile			
Start Address	start_addr			
End Address	end_addr			
MAC Filter Rules	mac_filter_rules			
MS AD User Data	ms_ad_user_data			
Next Server	nextserver			
Server Association Type	server_association_type	contains, equals, begins with, ends with		
Event Type	event_type	equals	RPZ LEASE TUNNEL NETWORK_IPV4 NETWORK_IPV6 RANGE_IPV4 RANGE_IPV6 FIXED_ADDRESS_IPV4 FIXED_ADDRESS_IPV6 HOST_ADDRESS_IPV6 HOST_ADDRESS_IPV6 SESSION	SESSION is used for the login and logout events for the session management templates.

Variables for DB Object Change Event - DHCP Fixed Address IPv4

*Variables for DB Object Change Event - DHCP Fixed Address IPv4*

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Timestamp	timestamp (ISO 8601 format)			The timestamp when the event occurs.
Infoblox Member IP	member_ip			Infoblox member IP (VIP or LAN1) that has generated the event.

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Infoblox Member Name	member_name			
WAPI object reference	_ref			
Comment	comment			
Disable	disable	equals		Boolean
Extensible Attributes	extattrs			Dictionary of extensible attributes
Network	network	equals, contained in		ip_addr/cidr
network_view	network_view	contains, equals, begins with, ends with		String format
MS AD User Data	ms_ad_user_data			
Name	name	contains, equals, begins with, ends with		
MAC Address	mac	contains, equals, begins with, ends with		
IPv4 Address	ipv4addr	equals, matches range, matches CIDR		
Event Type	event_type	equals	RPZ LEASE TUNNEL NETWORK_IPV4 NETWORK_IPV6 RANGE_IPV4 RANGE_IPV6 FIXED_ADDRESS_IPV4 FIXED_ADDRESS_IPV6 HOST_ADDRESS_IPV4 HOST_ADDRESS_IPV6 SESSION	SESSION is used for the login and logout events for the session management templates.

Variables for DB Object Change Event - DHCP Fixed Address IPv6

*Variables for DB Object Change Event - DHCP Fixed Address IPv6*

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Timestamp	timestamp (ISO 8601 format)			The timestamp when the event occurs.
Infoblox Member IP	member_ip			Infoblox member IP (VIP or LAN1) that has generated the event.
Infoblox Member Name	member_name			
WAPI object reference	_ref			
Comment	comment			
Disable	disable	equals		Boolean
Extensible Attributes	extattrs			Dictionary of extensible attributes
Network	network	equals, contained in		ip_addr/cidr
network_view	network_view	contains, equals, begins with, ends with		String format
MS AD User Data	ms_ad_user_data			
Name	name	contains, equals, begins with, ends with		
IPv6 DUID	duid	contains, equals, begins with, ends with		
Address Type	address_type	equals		Address, Prefix, or Both
IPv6 Address	ipv6addr	equals, matches range, matches CIDR		
IPv6 Address Prefix	ipv6prefix	equals, matches range, matches CIDR		
IPv6 Address Prefix bits	ipv6prefix_bits	equals		

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Event Type	event_type	equals	RPZ LEASE TUNNEL NETWORK_IPV4 NETWORK_IPV6 RANGE_IPV4 RANGE_IPV6 FIXED_ADDRESS_IPV4 FIXED_ADDRESS_IPV6 HOST_ADDRESS_IPV4 HOST_ADDRESS_IPV6 SESSION	SESSION is used for the login and logout events for the session management templates.

#### Variables for DB Object Change Event - DHCP Host Address IPv4

##### *Variables for DB Object Change Event - DHCP Host Address IPv4*

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Timestamp	timestamp (ISO 8601 format)			The timestamp when the event occurs.
Infoblox Member IP	member_ip			Infoblox member IP (VIP or LAN1) that has generated the event.
Infoblox Member Name	member_name			
WAPI object reference	_ref			
Extensible Attributes	extattrs			Dictionary of extensible attributes from parent host record
Network	network	equals, contained in		ip_addr/cidr
network_view	network_view	contains, equals, begins with, ends with		String format
MS AD User Data	ms_ad_user_data			
Host	host	contains, equals, begins with, ends with		

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
MAC Address	mac	contains, equals, begins with, ends with		
IPv4 Address	ipv4addr	equals, matches range, matches CIDR		
Event Type	event_type	equals	RPZ LEASE TUNNEL NETWORK_IPV4 NETWORK_IPV6 RANGE_IPV4 RANGE_IPV6 FIXED_ADDRESS_IPV4 FIXED_ADDRESS_IPV6 HOST_ADDRESS_IPV4 HOST_ADDRESS_IPV6 SESSION	SESSION is used for the login and logout events for the session management templates.

Variables for DB Object Change Event - DHCP Host Address IPv6

*Variables for DB Object Change Event - DHCP Host Address IPv6*

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Timestamp	timestamp (ISO 8601 format)			The timestamp when the event occurs.
Infoblox Member IP	member_ip			Infoblox member IP (VIP or LAN1) that has generated the event.
Infoblox Member Name	member_name			
WAPI object reference	_ref			
Extensible Attributes	extattrs			Dictionary of extensible attributes from parent host record

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Network	network	equals, contained in		ip_addr/cidr
network_view	network_view	contains, equals, begins with, ends with		String format
MS AD User Data	ms_ad_user_data	contains, equals, begins with, ends with		
Host	host	contains, equals, begins with, ends with		
IPv6 DUID	duid	contains, equals, begins with, ends with		
Address Type	address_type	equals		Address, Prefix, or Both
IPv6 Address	ipv6addr	equals, matches range, matches CIDR		
IPv6 Address Prefix	ipv6prefix	equals, matches range, matches CIDR		
IPv6 Address Prefix bits	ipv6prefix_bits	equals		
Event Type	event_type	equals	RPZ LEASE TUNNEL NETWORK_IPV4 NETWORK_IPV6 RANGE_IPV4 RANGE_IPV6 FIXED_ADDRESS_IPV4 FIXED_ADDRESS_IPV6 HOST_ADDRESS_IPV4 HOST_ADDRESS_IPV6 SESSION	SESSION is used for the login and logout events for the session management templates.



### Variables for DNS Record

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Record Name	record_name	contains, equals, begins with, ends with		
Record Type	record_type	equals		A Record, AAAA Record, CAA Record, NS Record, MX Record, Alias Record, PTR Record, NAPTR Record, CNAME Record, DNAME Record, TLSA Record, TXT Record, SOA Record, SRV record, Unknown Record
Auto created Records	auto_rec			
Network View	network_view	contains, equals, begins with, ends with		
Zone Name	zone_name	contains, equals, begins with, ends with		
DNS View	dns_view	contains, equals, begins with, ends with		
USER NAME	USER_NAME			Filtered by user name

### Variables for DNS Zone

NIOS Field Name	Template Variable Name	Supported Filter	Enriched Data	Comment
Zone Name	zone_name	contains, equals, begins with, ends with		
Zone Type	zone_type	equals		Authoritative Zone, Delegated Zone, Forward Zone, Response Policy Zone, Stub Zone
Network View	network_view	contains, equals, begins with, ends with		
DNS View	dns_view	contains, equals, begins with, ends with		
USER NAME	USER_NAME			Filtered by user name

## Result Parsing

Besides the R: name space, which is always initialized, if a template step includes the parse setting, then the result returned by the server is interpreted and will be available in the P: name space. Result parsing is extremely naive and supports only the following:

- **JSON:** In this case, the returned JSON will be available in the P name space if it is a dictionary, otherwise it will be as-is in the P:PARSE variable. If JSON is malformed, an error will be logged and P:PARSE will be set to an empty dictionary.
- **REGEX:** The supplied REGEX is applied to the whole body and it will be available in P:PARSE[0] to P:PARSE[N], if you have specified more than one groupings. In this case, P:PARSE[0] contains the full match, and P:PARSE[1..N] contains each individual grouping match. If no groupings are parsed, then P:PARSE will be a string with the matched expression. If one grouping is parsed, then P:PARSE will be a string with the matched grouping result. Note that ^ and \$ anchors are the anchors for the whole output to be parsed.
- **REGEXLINE:** The supplied REGEX is applied to every line returned by the server, and each match is assigned to P:PARSE[0] to P:PARSE[N] depending on how many lines match. The REGEX must contain only one grouping. If there are many groupings then the last matched grouping is put in the PARSE value, but if there are no groupings provided then the full match is put in each line. Note that ^ and \$ anchors in this case are the anchors for each line.
- **REGEXMULTILINE:** REGEXMULTILINE can be used when multiple groupings are required to be matched in multiple lines. Each match will be in P:PARSE[0..n] and each individual match is a list where the first value is the full REGEX match, and each subsequent value is a grouping match. The first grouping match of the first match can be accessed via \${P:PARSE[0][1]}. Note that ^ and \$ anchors in this case are the anchors for each line.

For all REGEX cases, if there are no matches, P:PARSE is set to an empty string.

- **CONDITION:** In some cases, when a condition on a REGEX match is required, it is recommended to use a CONDITION step with the :L (length) format specifier applied to PARSE. So when \${P:L:PARSE} is matched with = 0, it would create a condition evaluating to true if there was a regular expression match.
- **XML:** For XML, the XML data is converted into a dictionary of dictionaries/lists depending on the XML present. Similar to JSON, this will be available in the P name space if it is a dictionary, otherwise it will be as-is in the P:PARSE variable. If the XML is malformed or not parsable, then an error is logged and P:PARSE is set to an empty dictionary.

Parsing does not support DTDs, schemas, or XML attributes. It simply converts the XML document as-is. This also means that if the schema defines a particular element to be a list, it might not be deserialized as a list depending on how many members are present (if there is only one, then the parent is not considered a list.) When serializing XML, any variable composed of a dictionary with an ' <xmla> ' member set to True will be serialized as an XML element with attributes.

If an XMLA element is serialized under JSON/XML, the attributes will be ignored and the appliance returns a meaningful result (a list of XMLA variables will be serialized in JSON as a list of the values of these elements, assuming the values are simple strings/numbers).

For example, consider the following XML:

```
<?xml version="1.0" encoding="UTF-8" ?>
<!DOCTYPE APPLIANCE_LIST_OUTPUT SYSTEM
"https://qualysapi.qualys.com/api/2.0/fo/appliance/appliance_list_
output.dtd">
<APPLIANCE_LIST_OUTPUT>
  <RESPONSE>
    <DATETIME>2014-01-02T09:26:01Z</DATETIME>
    <APPLIANCE_LIST>
      <APPLIANCE>
```

```

<ID>777</ID>
<NAME>scanner1</NAME>
<SOFTWARE_VERSION>2.6</SOFTWARE_VERSION>
<RUNNING_SCAN_COUNT>0</RUNNING_SCAN_COUNT>
<STATUS>Online</STATUS><S2>Online</S2>
</APPLIANCE>
<APPLIANCE>
<ID>1127</ID>
<NAME>scanner2</NAME>
<SOFTWARE_VERSION>2.6</SOFTWARE_VERSION>
<RUNNING_SCAN_COUNT>0</RUNNING_SCAN_COUNT>
<STATUS>Online</STATUS><S2>Online</S2>
</APPLIANCE>
<APPLIANCE>
<ID>1131</ID>
<NAME>scanner3</NAME>
<SOFTWARE_VERSION>2.6</SOFTWARE_VERSION>
<RUNNING_SCAN_COUNT>0</RUNNING_SCAN_COUNT>
<STATUS>Offline</STATUS><S2>Online</S2>
</APPLIANCE>
</APPLIANCE_LIST>
<LICENSE_INFO>
<QVSA_LICENSES_COUNT>10</QVSA_LICENSES_COUNT>
<QVSA_LICENSES_USED>3</QVSA_LICENSES_USED>
</LICENSE_INFO>
</RESPONSE>
</APPLIANCE_LIST_OUTPUT>

```

The XML will be deserialized as follows:

```

{
  "APPLIANCE_LIST_OUTPUT":
  {
    "RESPONSE": {
      "APPLIANCE_LIST": [
        {

```

```
        "ID": "777",
        "NAME": "scanner1",
        "RUNNING_SCAN_COUNT": "0",
        "S2": "Online",
        "SOFTWARE_VERSION": "2.6",
        "STATUS": "Online"
    },
    {
        "ID": "1127",
        "NAME": "scanner2",
        "RUNNING_SCAN_COUNT": "0",
        "S2": "Online",
        "SOFTWARE_VERSION": "2.6",
        "STATUS": "Online"
    },
    {
        "ID": "1131",
        "NAME": "scanner3",
        "RUNNING_SCAN_COUNT": "0",
        "S2": "Online",
        "SOFTWARE_VERSION": "2.6",
        "STATUS": "Offline"
    }
],
"DATETIME": "2014-01-02T09:26:01Z",
"LICENSE_INFO": {
    "QVSA_LICENSES_COUNT": "10",
    "QVSA_LICENSES_USED": "3"
}
```

```
    },  
  }  
}
```

But the following will be deserialized differently:

```
<?xml version="1.0" encoding="UTF-8" ?>  
<!DOCTYPE APPLIANCE_LIST_OUTPUT SYSTEM  
"https://qualysapi.qualys.com/api/2.0/fo/appliance/  
appliance_list_output.dtd">  
<APPLIANCE_LIST_OUTPUT>  
  <RESPONSE>  
    <DATETIME>2014-01-02T09:26:01Z</DATETIME>  
    <APPLIANCE_LIST>  
      <APPLIANCE>  
        <ID>777</ID>  
        <NAME>scanner1</NAME>  
        <SOFTWARE_VERSION>2.6</SOFTWARE_VERSION>  
        <RUNNING_SCAN_COUNT>0</RUNNING_SCAN_COUNT>  
        <STATUS>Online</STATUS><S2>Online</S2>  
      </APPLIANCE>  
    </APPLIANCE_LIST>  
    <LICENSE_INFO>  
      <QVSA_LICENSES_COUNT>10</QVSA_LICENSES_COUNT>  
      <QVSA_LICENSES_USED>3</QVSA_LICENSES_USED>  
    </LICENSE_INFO>  
  </RESPONSE>  
</APPLIANCE_LIST_OUTPUT>
```

The XML will be deserialized as follows (note the difference in the `appliance_list`):

```
{  
  "APPLIANCE_LIST_OUTPUT":
```

```

    {
      "RESPONSE": {
        "APPLIANCE_LIST": {
          "APPLIANCE": {
            {
              "ID": "777",
              "NAME": "scanner1",
              "RUNNING_SCAN_COUNT": "0",
              "S2": "Online",
              "SOFTWARE_VERSION": "2.6",
              "STATUS": "Online"
            }
          }
        },
        "DATETIME": "2014-01-02T09:26:01Z",
        "LICENSE_INFO": {
          "QVSA_LICENSES_COUNT": "10",
          "QVSA_LICENSES_USED": "3"
        }
      }
    }
  }
}

```

- **XMLA:** Infoblox strongly recommends that you use XMLA as the quoting option (as opposed to XML) when you create new action templates. New operations such as PUSH, POP, SHIFT and others are not officially supported for XML parsed data.

XMLA parsing strips all white spaces (horizontal tab, line feed, vertical tab, form feed, carriage return, and space). To keep all white spaces, XMLA\_WHITESPACE enum is added. The parsing is the same as XMLA except that there is no white space stripping.

If you set parse to XMLA, the XML parsing supports XML attributes, and the XML document is put in the P:PARSE namespace variable, as illustrated in the following example. This example contains a mixed-attribute XML message with some values that contain attributes and some do not.

```

<?xml version="1.0" encoding="utf-8"?>
<SiteConfigResponse success="1">
  <Site id="27" name="SOAPUI13006925d-7dac-428d-

```

```

aaf1-4038a98838a1" description="" riskfactor="1.0" isDynamic="0">
  <Description/>
  <Hosts a0="123">
    <host a1="123">server1.example.com</host>
    <host>server2.example.com</host>
    <host>server3.example.com</host>
    <host>server4.example.com</host>
    <host>server5.example.com</host>
  </Hosts>
  <Credentials></Credentials>
  <Alerting>
    <Alert name="test" enabled="1" maxAlerts="2">
      <scanFilter scanStart="1" scanStop="1" scanFailed="1"
scanResumed="1" scanPaused="1"/>
      <vulnFilter severityThreshold="1" confirmed="1"
unconfirmed="1" potential="1"/>
      <smtpAlert sender="user1@example.com" server="server6.example.com"
limitText="0">
        <recipient>user2@example.com</recipient>
      </smtpAlert>
    </Alert>
  </Alerting>
  <ScanConfig configID="28" name="Full audit" templateID="full-audit"
engineID="3" configVersion="3">
    <Schedules></Schedules>
  </ScanConfig>
</Site>
</SiteConfigResponse>

```

The following is the deserialized response in P:PARSE when using XMLA parsing:

```

[
  {
    "<xmla>": true,
    "attrs": {
      "success": "1"
    }
  },

```

```
"index": {
  "Site": 0
},
"name": "SiteConfigResponse",
"value": [
  {
    "<xmla>": true,
    "attrs": {
      "description": "",
      "id": "27",
      "isDynamic": "0",
      "name": "SOAPUI13006925d-7dac-428d-aaf1-4038a98838a1",
      "riskfactor": "1.0"
    },
    "index": {
      "Alerting": 3,
      "Credentials": 2,
      "Description": 0,
      "Hosts": 1,
      "ScanConfig": 4
    },
    "name": "Site",
    "value": [
      {
        "<xmla>": true,
        "attrs": {},
        "index": {},
        "name": "Description",
        "value": []
      },
      {
        "<xmla>": true,
        "attrs": {
          "a0": "123"
        }
      }
    ]
  }
]
```



```
    },
    "index": {
      "host": 4
    },
    "name": "Hosts",
    "value": [
      {
        "<xmla>": true,
        "attrs": {
          "a1": "123"
        },
        "index": {},
        "name": "host",
        "value": ["server1.example.com"]
      },
      {
        "<xmla>": true,
        "attrs": {},
        "index": {},
        "name": "host",
        "value": ["server2.example.com"]
      },
      {
        "<xmla>": true,
        "attrs": {},
        "index": {},
        "name": "host",
        "value": ["server3.example.com"]
      },
      {
        "<xmla>": true,
        "attrs": {},
        "index": {},
        "name": "host",
        "value": ["server4.example.com"]
      }
    ]
  },
}
```

```
{
  "<xmla>": true,
  "attrs": {},
  "index": {},
  "name": "host",
  "value": ["server5.example.com"]
}
],
{
  "<xmla>": true,
  "attrs": {},
  "index": {},
  "name": "Credentials",
  "value": []
},
{
  "<xmla>": true,
  "attrs": {},
  "index": {
    "Alert": 0
  },
  "name": "Alerting",
  "value": [
    {
      "<xmla>": true,
      "attrs": {
        "enabled": "1",
        "maxAlerts": "2",
        "name": "test"
      },
      "index": {
        "scanFilter": 0,
        "smtpAlert": 2,
        "vulnFilter": 1
      },
      "name": "Alert",
```

```
        "value": [
          {
            "<xmla>": true,
            "attrs": {
              "scanFailed": "1",
              "scanPaused": "1",
              "scanResumed": "1",
              "scanStart": "1",
              "scanStop": "1"
            },
            "index": {},
            "name": "scanFilter",
            "value": []
          },
          {
            "<xmla>": true,
            "attrs": {
              "confirmed": "1",
              "potential": "1",
              "severityThreshold": "1",
              "unconfirmed": "1"
            },
            "index": {},
            "name": "vulnFilter",
            "value": []
          },
          {
            "<xmla>": true,
            "attrs": {
              "limitText": "0",
              "sender": "user1@example.com",
              "server": "server6.example.com"
            },
```

```

        "index": {
            "recipient": 0
        },
        "name": "smtpAlert",
        "value": [
            {
                "<xmla>": true,
                "attrs": {},
                "index": {},
                "name": "recipient",
                "value":
["user2@example.com"]
            }
        ]
    }
]
},
{
    "<xmla>": true,
    "attrs": {
        "configID": "28",
        "configVersion": "3",
        "engineID": "3",
        "name": "Full audit",
        "templateID": "full-audit"
    },
    "index": {
        "Schedules": 0
    },
    "name": "ScanConfig",
    "value": [
        {
            "<xmla>": true,
            "attrs": {},
            "index": {},

```

```

        "name": "Schedules",
        "value": []
    }
]

```

As shown in the example above, any XML value will become an internal dictionary that contains separate ' `attrs` ' and ' `value` ' members representing the XML attributes (if exist) and the XML values (if exist) of the element respectively, as well as the single "" boolean (set to `True` ) to qualify this particular field as an XML attribute field (this is used when serializing, as well as for user-created XMLA values, see below).

You typically address the XMLA variable by using `VAR{tag}{subtag}` , which uses the VALUES of the tag (its value or its subtag(s)). To access the attributes of a tag, you use `VAR{tag}{{attributename}}` .

Example: `P::PARSE{SiteConfigResponse}{{success}}`

If you need to access the name of a tag instead (for example if the remote server can return different tags depending on the status), use the `[[name]]` syntax. For example, you can use `P::PARSE[[name]]` .

### Repeating a Parse Operation

You can use the `XC:PARSE` operation to parse the body once again using another method. `R:BODY` contains the body of the response.

#### Example:

Having the L namespace as:

```

{
  "config": {
    "parse": "JSON"
  },
}

```

`L:data` contains the result of the parsing operation after the following step:

```

{
  "version": "4.0",
  ...
  "steps": [
    ...
    {
      "name": "repeat_parse",
      "operation": "NOP",
      "body_list": [
        "${XC:COPY:{L:data}:{R:BODY}}",
        "${XC:PARSE:{L:config}:{L:data}}"
      ]
    },
    ...
  ],
  ...
}

```

## Configuring Outbound Endpoints

An endpoint sends outbound notifications based on the notification rule and the outbound template that you have configured. With NIOS, you can configure REST API, DXL, Syslog, Cisco ISE endpoints to send outbound notifications. You can use the RESTful API and DXL fabric to obtain core network service information from the Infoblox Grid to assist with profiling the source or destination of network devices or use the RESTful API and WAPI in DXL endpoint to change configurations in the Infoblox Grid to help mitigate security threats. In addition to querying inbound data and changing system configurations and query interfaces, you can use the RESTful API and DXL messages to send outbound notifications so you can prioritize your security needs by detecting new hosts or networks or managing network access control.

The REST API endpoint you configure must be REST enabled so that they can handle RESTful API calls. The DXL endpoints must be connected to DXL brokers and listen on specific DXL topics as configured in the DXL action template. You must upload session management and action templates before you configure endpoints.



### Note

Infoblox recommends that you send notifications from a Grid Master Candidate, when it is available, instead of the Grid Master.

## Configuring REST API Endpoints

You can configure REST API endpoints and define rules to send outbound notifications to the REST enabled target system.

To configure a REST API endpoint, complete the following:

1. From the **Grid/System** tab, select the **Ecosystem** tab -> **Outbound Endpoint** tab and then click **Add** -> **Add REST API Endpoint** from the Toolbar.
2. In the **Add REST API Endpoint** wizard, complete the following:
  - **URI**: Specify the URI for the endpoint to which you are sending the outbound notifications. Example: `https://10.36.101.14/offices`.
  - **Test Connection**: Click this to validate the endpoint settings and test the connectivity between the Grid Master and the endpoint. It also tests the connection between the Grid Master Candidate that is assigned as the outbound member and the endpoint. Grid Manager displays a message indicating whether the connection is successful. Note that the test does not validate username, password, or certificate for the endpoint. It only tests the basic connection between the Grid Master and the endpoint.
  - **Name**: Specify the name used to identify the endpoint.
  - **Vendor Type**: The REST API vendor type associated with the endpoint. This is optional.
  - **Network View**: This appears only when you have multiple network views. From the drop-down list, select the network view in which you want to create the network.
  - **Auth Username**: Enter the username of the target endpoint. The appliance ignores the **Auth Username** for WAPI related steps in any action templates if WAPI integration is configured. It still uses this username for other non-WAPI related steps.
  - **Auth Password**: Enter the user account password for the target endpoint. You can click **Clear Password** to clear the password and set a new one. The appliance ignores the **Auth Username** for WAPI related steps in any action templates if WAPI integration is configured. It still uses this password for other non-WAPI related steps.
  - **Client Certificate**: Click **Select** to upload the endpoint certificate. In the *Upload* dialog box, click **Select** to navigate to the certificate, and then click **Upload**.
  - **WAPI Integration Username**: If you have included at least one "wapi" related field in your action template, you must configure WAPI integration; otherwise the WAPI step will fail due to an authorization error. Enter the username of the admin user you want to designate for RESTful API outbound notifications. The appliance ignores the **Auth Username** and **Auth Password** for WAPI related steps in any action templates if WAPI integration is configured.
  - **WAPI Integration Password**: Enter the password of the admin user you have designated for RESTful API outbound notifications.
  - **Server Certificate Validation**: Select one of the following for server certificate validation:
    - **Use CA Certificate Validation (Recommended)**: Select this to validate the CA certificate for the endpoint. The certificate is used to establish a secure connection to the endpoint before data transmission. Click **CA Certificates** to upload the trusted CA certificate of the endpoint. In the *CA Certificates* dialog box, click the Add icon, and then navigate to the certificate to upload it. This is the default.
      - **Enable Host Validation**: Select this to enable the validation of the hostname for the endpoint, in addition to the CA certificate. If you do not select this, the appliance validates only the CA certificate.
      - **Do not use validation (Not recommended)**: Infoblox does not recommend using this for your production system. Use this for testing purposes only.
  - **Member Source outbound API requests from**: Select the one of the following to process and send outbound API notifications:
    - **Selected Grid Master Candidate (Recommended)**: Select this to use the Grid Master Candidate to process and send outbound notifications to the endpoint. If there are multiple Grid Master candidates, select a Grid Master Candidate from the drop-down list. This is the recommended choice and is selected by default because the CPU and memory required for processing and sending outbound events from the Grid Master Candidate can be offloaded or manually load balanced across multiple Grid Master Candidates if required.
    - **Current Grid Master**: Click this to use the Grid Master to send outbound notifications to the endpoint. When you use the Grid Master as the outbound member, ensure that it has enough CPU and memory to process all the workloads and processes, in addition to being an outbound

member. Infoblox recommends that you use the Grid Master as an outbound member only for testing purposes to avoid overloading the Grid Master and to maintain optimal performance for the Grid.

- **Comment:** Enter additional information about the REST API endpoint.
  - **Disable:** Select this if you want to save the configuration but do not want to use it yet. You can clear this checkbox when you are ready to use this configuration.
3. Click **Next** to set the duration of time that the endpoint waits for a response from the outbound member. Complete the following to specify session timeout value:
    - **Timeout:** Specify the session timeout value for the endpoint. The default value is 30 seconds.
    - **Log Level:** From the drop-down list, select the severity level for the events. The severity level you select here determines the type of events that are being logged. This can be **Debug**, **Info**, **Warning**, or **Error**. When you select **Debug**, all fields or variables used in the events that were sent to the endpoint are logged, including deduplicated events for RPZ hits. For information about deduplication, see [Deduplicating Events](#). Note that setting this to **Debug** might slightly affect the performance of your production system.
    - **Template:** Click **Select Template** to select a session management template. For information, see [Creating Session Management Templates](#).
    - **Vendor Type:** Displays the vendor information for the endpoint.
    - **Template Type:** Displays the Session Management template.
    - **Parameters:** Displays the parameters of the template you select. You can access these values in the notification rules.
  4. Click **Next** to add extensible attributes for the endpoint. For information, see [Managing Extensible Attributes](#).
  5. Save the configuration.



#### Note

If your outbound member is a Grid Master Candidate and in case the Grid Master Candidate is promoted to the Grid Master, make sure that you modify the outbound member to the Grid Master on the endpoint configuration to avoid any outbound notification failures.

## Configuring DXL Endpoints

When adding a DXL endpoint, you must configure the DXL client and the broker. DXL brokers that are installed on managed systems route messages between connected clients. The network of brokers tracks active consumers and dynamically adjusts the message routing as needed. As shown in the figure below, a broker relays a message when a client requests a service or when an update is broadcast.



An Outbound worker that acts as a DXL client sends data and templates using the DXL protocol to the DXL brokers fabric. You can change the format of the DXL message using relevant template. A connection is established as soon as the outbound worker starts transmitting the data.

You can manually configure the list of DXL brokers that are used by DXL clients in NIOS or import the broker configuration file. The DXL endpoint configuration requires import of DXL brokers list and their certificates on the Infoblox side as well as import of Infoblox certificate on the McAfee side. Note that you must install the Security Ecosystem license before you configure a DXL endpoint.

For a McAfee ePolice Orchestrator, you must do the following:

1. Import NIOS certificate.
2. Export DXL broker certificates.
3. Export a DXL broker list.

To configure a DXL endpoint, complete the following:



1. From the **Grid/System** tab, select the **Ecosystem** tab -> **Outbound Endpoint** tab and then click **Add** -> **Add DXL Endpoint** from the Toolbar.
2. In the **Add DXL Endpoint** wizard, complete the following:
  - **Name:** Specify the name used to identify the endpoint.
  - **Vendor Type:** The DXL vendor type associated with the endpoint. This is optional.
  - **Client Certificate:** Click **Generate** to generate and upload both the client and CA certificates of the endpoint on NIOS. When you click **Generate**, the client certificate is automatically uploaded on NIOS and a copy of CA certificate is downloaded. Import this downloaded CA certificate to the DXL server. For information about how to import the CA certificate, refer to the McAfee documentation. If you already have the client certificate, you can upload it by clicking the Upload icon. Click **Upload** to upload the client certificate. In the *Upload* dialog box, click **Select** to navigate to the certificate, and then click **Upload**.
  - **CA Certificates:** Click **CA Certificates** to upload the broker Certificate. Download the broker certificate from the DXL server and upload it to NIOS. In the *CA Certificates* dialog box, click the Add icon, and then navigate to the certificate to upload it.
  - **WAPI Integration Username:** If you have included at least one “wapi” related field in your action template, you must configure WAPI integration; otherwise, the WAPI step will fail due to an authorization error. Enter the username of the admin user you want to designate for DXL notifications.
  - **WAPI Integration Password:** Enter the password of the admin user you have designated for DXL notifications.
  - **Member Source outbound API requests from:** Select one of the following to process for sending outbound notifications:
    - **Selected Grid Master Candidate (Recommended):** Select this to use the Grid Master Candidate to process and send outbound notifications to the endpoint. If there are multiple Grid Master candidates, select a Grid Master Candidate from the drop-down list. This is the recommended choice and is selected by default because the CPU and memory required for processing and sending outbound events from the Grid Master Candidate can be offloaded or manually load balanced across multiple Grid Master Candidates if required.
    - **Current Grid Master:** Click this to use the Grid Master to send outbound notifications to the endpoint. When you use the Grid Master as the outbound member, ensure that it has enough CPU and memory to process all the workloads and processes, in addition to being an outbound member. Infoblox recommends that you use the Grid Master as an outbound member only for testing purposes to avoid overloading the Grid Master and to maintain optimal performance for the Grid.
  - **Comment:** Enter additional information about the DXL endpoint.
  - **Disable:** Select this if you want to save the configuration but do not want to use it yet. You can clear this checkbox when you are ready to use this configuration.
3. Click **Next** to add the DXL broker. There are two ways to configure the DXL broker. You can either manually enter the host name of the broker or you can import the broker configuration file using the **Import** option.
 

**To create your own file with brokers list:**  
In the *Brokers* wizard, complete the following:

  - Click **Add** to open the *Add Broker* wizard. Enter the host name in the Host Name text box. Optionally, you can enter the following information as well:
    - **IP address:** Enter the IP address of the DXL broker.
    - **Unique ID:** A unique identifier for the broker. This is useful for identifying the DXL broker in log messages.
    - **Port information:** The port number used to communicate with the DXL broker.
  - **To import the broker configuration file:**  
**In the *Brokers* wizard, complete the following:**
    - Click **Import** to upload the broker configuration file. In the *Upload* dialog box, click **Select** to navigate to the certificate. You can export the Broker configuration file `brokerlist.properties` file from McAfee ePolice Orchestrator (McAfee ePO). For information how to export, refer to the McAfee documentation. Click **Upload** to upload the broker configuration file.
    - Click **Test Connection** to validate the connectivity between the DXL broker fabric and the Grid Master.
4. Click **Next** to add a DXL topic. DXL uses topics to send data. You can then add the topic to a notification rule so that NIOS can send notifications when an event related to the topic occurs.
 

**To add a topic:**

- Click the **Topics** tab. Click the Add icon to enter a topic. Topics may be in the format defined in the session management template. For example, /infoblox/outbound/LEASE.
5. Click **Next** to set the severity level for the events.
    - **Timeout:** Specify the session timeout value for the endpoint. The default value is 30 seconds.
    - **Log Level:** From the drop-down list, select the severity level for the events. The severity level you select here determines the type of events that are being logged. This can be **Debug**, **Info**, **Warning**, or **Error**. When you select **Debug**, all fields or variables used in the events that were sent to the endpoint are logged, including deduplicated events for RPZ hits. For information about deduplication, see [Deduplicating Events](#). Note that setting this to **Debug** might slightly affect the performance of your production system.
    - **Vendor Type:** Displays the vendor information for the endpoint.
    - **Template Type:** Displays the Session Management template.
    - **Parameters:** Displays the parameters of the template you select. You can access these values in the notification rules.
  6. Click **Next** to add extensible attributes for the endpoint. For information, see [Managing Extensible Attributes](#).
  7. Save the configuration.



#### Note

If your outbound member is a Grid Master Candidate and in case the Grid Master Candidate is promoted to the Grid Master, make sure that you modify the outbound member to the Grid Master on the endpoint configuration to avoid any outbound notification failures.

## Configuring Syslog Endpoints

You can configure syslog endpoints to define syslog message format. When an event is triggered, the syslog message is sent based on the format you define. You can then analyze the data presented in the messages and take corrective measures. To do this, you must configure syslog endpoints. You can send syslog notifications either in raw or formatted text and also send a test syslog notification.

To configure a syslog endpoint, complete the following:

1. From the **Grid/System** tab, select the **Ecosystem** tab -> **Outbound Endpoint** tab and then click **Add** -> **Add Syslog Endpoint** from the Toolbar.
2. In the **Add Syslog Endpoint** wizard:
  - **Name:** Specify a name for the endpoint.
  - Click the + icon to add a syslog address:
  - **Address:** Enter the IP address of the syslog server.
  - **Transport:** Select the connection type that the syslog server will use. Supported types are UDP, TCP, and Secure TCP. If you select TCP or UDP, the default port number is 514 and you do not need to upload a certificate. If you select Secure TCP, the default port number is 6514 and you need to upload a certificate.
  - **Certificate:** If you selected Secure TCP, you must upload an HTTPS or a CA certificate. For more information, see [Managing Certificates](#).
  - **Port:** Specify the port number that the syslog server will use to communicate with NIOS.
  - **Message Format:** Select the format of the sys log message. If you select Formatted, you must specify the facility and severity to be sent in the syslog message header.
  - **Host Name:** If you selected Formatted as the message format, then the value that you select from the **Host Name** drop-down list is sent in the syslog message header.
  - **Facility:** Select the location that determines the processes and daemons from which the log messages are generated.
  - **Severity:** Select a severity for the syslog message. The severity type that you select is sent in the syslog message header.
  - Click **Add**. The syslog server details are added to the table. You can add more syslog addresses by clicking the + icon. You can also generate a test syslog notification by clicking **Test**.
  - **Vendor Type:** Select the vendor information for the endpoint.
  - **WAPI Integration Username:** If you have included at least one "wapi" related field in your action template, you must configure WAPI integration; otherwise the WAPI step fails due to an authorization error. Enter the user name of the admin user you want to designate for Syslog outbound notifications. The appliance

- ignores the **Auth Username** and **Auth Password** for WAPI related steps in any action templates if WAPI integration is configured.
- **WAPI Integration Password:** Enter the password of the admin user you have designated for Syslog outbound notifications.
  - **Member Source outbound API requests from:** Select the one of the following to process and send outbound notifications:
  - **Comment:** Enter additional information about the REST API endpoint.
  - **Disable:** Select this if you want to save the configuration but do not want to use it yet. You can clear this checkbox when you are ready to use this configuration.
3. Click **Next** to set the duration of time that the endpoint waits for a response from the outbound member. Complete the following to specify session timeout value:
    - a. **Timeout:** Specify the session timeout value for the endpoint. The default value is 30 seconds.
    - b. **Log Level:** From the drop-down list, select the severity level for the events. The severity level you select here determines the type of events that are being logged. This can be **Debug**, **Info**, **Warning**, or **Error**. When you select **Debug**, all fields or variables used in the events that were sent to the endpoint are logged, including deduplicated events for RPZ hits. Note that setting this to **Debug** might slightly affect the performance of your production system.
    - c. **Template:** Click **Select Template** to select a session management template.
    - d. **Vendor Type:** Displays the vendor information for the endpoint.
    - e. **Template Type:** Displays the Session Management template.
    - f. **Parameters:** Displays the parameters of the template you select. You can access these values in the notification rules.
  4. Click **Next** to add extensible attributes for the endpoint.
  5. Save the configuration.

### Configuring Cisco ISE Endpoints Using Outbound Endpoint

You can configure a Cisco ISE endpoint by either using the **Ecosystem > Cisco ISE Endpoint** tab or by using the **Ecosystem > Outbound Endpoint** tab. The instructions in this section pertain to configuring Cisco ISE 2.6, 2.7, 3.0, and 3.1 using the **Outbound Endpoint** tab.

You can configure Cisco ISE 3.1 only using the outbound endpoint.

To configure the Cisco ISE endpoint using the **Outbound Endpoint** tab, complete the following:

1. From the **Grid/System** tab, select the **Ecosystem** tab -> **Outbound Endpoint** tab and then click **Add** -> **Add Cisco ISE Endpoint** from the Toolbar.
2. In the **Add Cisco ISE Endpoint** wizard:
  - **Server Address:** Enter the IP address of the Cisco ISE.
  - **Name:** Specify a name for the endpoint.
  - **Subscribing Member:** Select a Grid Master Candidate that you want to subscribe as the client on the Cisco ISE. Or, you can select the current Grid Master as the subscribing member. This member interacts with the Cisco ISE to obtain contextual information for the subscribed data types.
  - **Vendor Type:** The vendor type associated with the endpoint. This is optional.
  - **Client Certificate:** Click **Select** to upload the client certificate. In the *Upload* dialog box, click **Select** to navigate to the certificate, and then click **Upload**.
  - **Manage Certificates:** Click **CA Certificates** to upload the self-signed certificate or CA certificate. In the *CA Certificates* dialog box, click the **Add** icon, and then navigate to the certificate to upload it.
  - **WAPI Integration Username:** If you have included at least one "wapi" related field in your action template, you must configure WAPI integration; otherwise the WAPI step fails due to an authorization error. Enter the user name of the admin user you want to designate for Cisco ISE outbound notifications. The appliance ignores the **Auth Username** and **Auth Password** for WAPI related steps in any action templates if WAPI integration is configured.
  - **WAPI Integration Password:** Enter the password of the admin user you have designated for Cisco ISE outbound notifications.
  - **Test Connection:** Click this to validate the endpoint settings and test the connectivity between the Grid Master and the endpoint. It also validates the certificate that you uploaded and tests the connection between the Grid Master Candidate that is assigned as the outbound member and the endpoint. Grid Manager displays a message indicating whether the connection is successful. Note that the test does not

- validate the user name and password for the endpoint. It only tests the basic connection between the Grid Master and the endpoint, and validates the certificate.
- **Comment:** Enter additional information about the Cisco ISE endpoint.
  - **Disable:** Select this checkbox if you want to save the configuration but do not want to use it yet. You can clear this checkbox when you are ready to use this configuration.
  - Click **Next** to set the duration of time that the endpoint waits for a response from the outbound member. Complete the following to specify session timeout value:
    - i. **Timeout:** Specify the session timeout value for the endpoint. The default value is 30 seconds.
    - ii. **Log Level:** From the drop-down list, select the severity level for the events. The severity level you select here determines the type of events that are being logged. This can be **Debug**, **Info**, **Warning**, or **Error**. When you select **Debug**, all fields or variables used in the events that were sent to the endpoint are logged, including deduplicated events for RPZ hits. Note that setting this to **Debug** might slightly affect the performance of your production system.
    - iii. **Template:** Click **Select Template** to select a session management template.
    - iv. **Vendor Type:** Displays the vendor information for the endpoint.
    - v. **Template Type:** Displays the Session Management template.
    - vi. **Parameters:** Displays the parameters of the template you select. You can access these values in the notification rules.
3. Click **Next** to specify the data types that you are interested to obtain from the Cisco ISE. The Cisco ISE shares information only for the subscribed data types. Complete the following to specify data types you want to collect from the Cisco ISE server:
    - **Subscription Settings:** Select the predefined data types to which you want to subscribe from the **Available Data Type** table. Use the arrows to move data types from the **Available Data Type** table to the **Selected Data Type** table. NIOS receives information for all data types in the **Selected Data Type** table.
    - **Map other data types to Extensible Attributes:** You can create extensible attributes and map these extensible attributes to receive additional Cisco ISE data values, such as IP address, MAC, NAS IP Address, NAS Port ID, EPS Status, Posture Status, Posture Timestamp, Endpoint Profile Name, Account Session ID, and Audit Session ID. Click the Add icon and map a Cisco ISE data type to an extensible attribute. You can also select a row and click the Delete icon to delete it.
  4. Click **Next** to add data types that you want to publish to the Cisco ISE server. Use the arrows to move data types from the **Available** table to the **Selected** table. NIOS publishes information only for the data types that are added in the **Selected** table.
  5. Click **Next** to add extensible attributes for the endpoint.
  6. Save the configuration.

## Modifying Outbound Endpoint Configuration

To modify an endpoint configuration:

1. From the **Grid/System** tab, select the **Ecosystem** tab -> **Outbound Endpoint** tab, click the **Action** icon next to the endpoint name and select **Edit** from the menu.
2. The <Endpoint Name> *Endpoint* editor provides the following tabs from which you can modify data:
  - **General:** You can modify the general information of an endpoint.
  - **Brokers:** You can modify the DXL broker configuration, as described in Configuring DXL Endpoints below. This tab is available only for DXL endpoints.
  - **Session Management:** You can edit the session timeout value and upload a new session management template.
  - **Extensible Attributes:** You can add, modify, and delete extensible attributes that are associated with an endpoint.
3. Save the configuration.

## Viewing All Outbound Endpoints


The **Outbound Endpoint** tab displays all outbound endpoints that are configured on the NIOS appliance.

To view the list of outbound endpoints:

1. From the **Grid/System** tab, select the **Ecosystem** tab, and click the **Outbound Endpoint** tab.
2. Grid Manager displays the following information for each endpoint:
  - **Name:** The name of the endpoint.

- **Endpoint Type:** The endpoint type, such as DXL or REST API.
- **URI:** The URI to which the outbound notifications are sent.
- **Vendor Type:** The vendor type associated with the endpoint.
- **Outbound Member:** The outbound member that processes and sends outbound notifications. This can be either the Grid Master Candidate or the Grid Master. Infoblox recommends that you select the Grid Master Candidate and this is selected by default.
- **Comment:** Additional information about the endpoint configuration.
- **Client Certificate Valid From:** The timestamp when the client certificate for a notification endpoint is created.
- **Client Certificate Valid To:** The timestamp when the client certificate for a notification endpoint expires.
- **Disabled:** Indicates whether the endpoint is disabled.
- **Site:** This is a predefined extensible attribute.

You can also do the following in this tab:

- Click the Action icon  and do the following:
  - **Edit:** Select this to modify the endpoint information.
  - **Delete:** Select this to delete an endpoint. Click **Yes** in the *Delete Confirmation (REST API Endpoint)* dialog box to delete an endpoint.
  - **View Debug Log:** Select this to view debug messages about all events associated with the selected endpoint. Through a separate browser, you can view the debug logs from all Grid members.
- Edit an outbound endpoint information.
  - Select the endpoint, and then click the Edit icon.
- Delete an outbound endpoint.
  - Select the endpoint, and then click the Delete icon.
- Export the list of outbound endpoints.
  - Click the Export icon.
- Print the list of outbound endpoints.
  - Click the Print icon.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria:
  - a. In the filter section, click **Show Filter** and define filter criteria for the quick filter.
  - b. Click **Save** and complete the configuration in the Save Quick Filter dialog box.

The appliance adds the quick filter to the quick filter drop-down list in the panel. Note that global filters are prefixed with [G], local filters with [L], and system filters with [S].

- Sort the outbound end points in ascending or descending order by column.

## Configuring Notification Rules

You can configure notification rules after you have uploaded outbound templates and configured outbound endpoints on the NIOS appliance. For information about adding outbound endpoints, see [Configuring Outbound Endpoints](#). To send outbound notifications from NIOS to the target endpoints, you must configure notification rules. When adding rules, you can select REST API, DXL, or Syslog endpoint and associate the correct action template to the rule. The appliance validates the event type specified in the template with the event type that you select in the notification rule. The parameters defined in a template decides the way NIOS specific data is presented to an endpoint. Each notification rule specifies the target endpoint, notification rule criteria, and the outbound template being used to take action for the matching events.

### Note

When you remove all notification rules associated with an endpoint, all the debug logs for that endpoint will also be removed.

While configuring notification rules, you can decide whether you want to reduce the number of redundant RPZ hits, ADP hits, and object change discovery data events. Oftentimes, these hits come from the same client IPs, query FQDNs, or networks. To avoid receiving excessive events at the endpoint, you can configure the appliance to remove or deduplicate

subsequent events (after sending the first event) within a certain time period. Depending on your configuration, the appliance sends the first event and deduplicates subsequent events that match your filtering criteria within the configured lookback interval. For more information, see [Deduplicating Events](#) below.

## Adding Notification Rules

To add notification rules, complete the following steps:

1. From the **Grid/System** tab, select the **Ecosystem** tab -> **Notification** tab, and then click the Add icon.  
or  
From the **Grid/System** tab, select the **Ecosystem** tab, and click **Add Notification Rule** in the Toolbar.
2. In the *Add Notification* wizard, complete the following.
  - **Name:** Enter the name of the notification rule.
  - **Target:** Click **Select Endpoint** to select the endpoint type. If there are multiple endpoints, the *All Endpoints Selector* dialog box is displayed, from which you can select an endpoint name, such as **Cisco ISE**.
  - **Target Type:** Displays the target type. You cannot change this.
  - **Comment:** Enter useful information about the notification rule.
  - **Disable:** Select this option to disable the notification rule.
3. Click **Next** and complete the following to configure notification rules for the selected endpoint:
  - **Event:** Depending on the licenses you have installed in the Grid, you can select the event types you want to apply to the notification rules. The outbound member collects data for the selected events based on your configuration. Note that if there is a significant amount of data or if the network bandwidth is not sufficient, the outbound member might drop some of the events. In this case, you can access the syslog to view the messages related to the dropped events. In addition to basic information (such as timestamp, member IP, network, and others), data collected for some event type might include enriched data, such as discovered data, parent network information, and associated extensible attributes.  
From the drop-down list, select the event types that you want to monitor for the notification rules:
    - **DNS RPZ:** Select this to collect data for RPZ events. The **DNS RPZ** event type is available only if you have installed the **RPZ** license in the Grid. When you select this event type, you can enable event deduplication in the next step so that the appliance can avoid sending excessive events to the endpoint based on your configuration.
    - **Object Change DNS Record:** Select this to collect data for DNS records. That is, if a DNS record is added, updated, or deleted, the notification rule is triggered and the event notification is sent to the target endpoint. Dynamic records are not supported.
    - **Object Change DNS Zone:** Select this to collect data for DNS zones. That is, if a zone is created, updated, or deleted, the notification rule is triggered and the event notification is sent to the target endpoint.
    - **DNS Tunneling:** Select this to collect data for DNS tunneling events.
    - **DHCP Leases:** Select this to collect data for DHCP leases. Since the same IP addresses might be used by multiple systems, the appliance matches both the IP and the MAC address or the DUID to ensure that the discovered data is most likely to be correct.
    - **DXL Events:** Select this to collect data from the topic to which you subscribed when configuring the DXL endpoint. For more information, see [Configuring DXL Endpoints](#).
    - **IPAM:** Select this to send IPAM data. No notification rule is required for this event type. For more information, see [Publishing Data](#).
    - **Security ADP:** Select this to collect data for threat protection events. You can specify the maximum domain level for query FQDN for outbound threat protection events. For more information, see [Enabling Query FQDN for Outbound Notifications](#) below. When you create outbound notifications for security ADP event types, the server collects event statistics every 15 seconds to avoid excessive threat protection events. Note that you can execute test rules in JSON format for Security ADP event types. For more information, see [Deduplicating Events](#) below.
    - **Object Change DHCP Fixed Address IPv4, Object Change DHCP Fixed Address IPv6, Object Change Network IPv4, Object Change Network IPv6, Object Change Range IPv4, Object Change Range IPv6, Object Change Host Address IPv4, Object Change Host Address IPv6, Object Change Discovery Data:** Select any of these event types to collect data for database changes in fixed addresses, DHCP ranges, networks, DNS host addresses, and discovery data. If you select **Object Change Discovery Data**, when unmanaged IP addresses or devices are created,

updated, or deleted, the notification rule is triggered and the event notification is sent to the target endpoint.

- **Schedule:** Select this to schedule the notification rule configuration. You can set up schedules on an hourly, daily, weekly, or on a monthly basis. You can also choose to schedule the event to occur only once. You cannot specify other event types when you select **Schedule** from the drop-down list. Note that you can execute test rules in JSON format when you schedule the notification rule configuration. You cannot choose an action rule when you schedule the notification rule configuration.
  - **Priority:** This field is displayed only if you select **Schedule** from the drop-down list. Select a priority value, **Normal** or **High**, for scheduled events from the drop-down list. When you select **Normal**, the event is added to the queue right after the existing events in the list and is executed after all events that are already scheduled. Select **High** if you want the event type to be executed soon after the execution of the current event in the list of events that are scheduled. For more information, see [Scheduling Tasks](#).
- **Action:** This field is displayed only if you have selected **Cisco ISE** as the endpoint (the Target field). Otherwise, this field is hidden.

In the **Match the following rule** section, select the filters, operators, and values from the drop-down lists for the selected event type. You can use the + icon to construct nested expressions for the rule. The filters change depending on what you selected as the event type. Some of the filters are:

- **DNS RPZ:** Action Policy, BloxOne Threat Defense Cloud Hit Class, BloxOne Threat Defense Cloud Hit Property, BloxOne Threat Defense Cloud Hit Type, Query Name, RPZ Name, RPZ Type, Rule Name, Source IP, and Threat Origin.
- **DNS Tunneling:** Source IP.
- **DHCP Leases:** DHCP Fingerprint and Lease State.
- **DXL Events:** DXL topic that you entered in the **Topics** field when configuring the DXL endpoint. For more information, see [Configuring DXL Endpoints](#).
- **Record Name:** Name of the DNS record. For example, AA, CNAME, SRV, and so on.
- **DNS Records:** Supported records are A, AAAA, AAA, CNAME, SRV, ALIAS, NS, PTR, MX, TXT, TLSA, CAA, SOA, DNAME, NAPTR, and UNKNOWN.
- **Zone Type:** Supported zones are Authoritative, Forward, Stub, Delegation, and RPZ.
- **User Name:** Name entered in the WAPI Integration Username field. For more information, see [Configuring Outbound Endpoints](#).
- **IPAM:** In the *Notify the target* section, there are predefined data types in the **Available** table you can publish. Click **Override** and use the arrows to move data types from the **Available** table to the **Selected** table and vice versa. The appliance sends information for all data types that are added to the **Selected** table. If you do not override, the publication settings are inherited from those configured while adding the Cisco ISE server. Note that you can configure only one IPAM rule per Cisco ISE server. For more information, see [Publishing Data](#).
- **Security ADP:** Rule Message, Hits Count, Member IP, Member Name, Query FQDN, Rule Action, Rule Category, Rule Severity, SID, and Source IP. When you select **Member Name**, the appliance displays all the ADP members that are available.
- **Object Change Fixed Address IPv4:** Disable, IPv4 Address, MAC, Name, Network, Network View, Username, and User Group.
- **Object Change Fixed Address IPv6:** Address Type, Disable, DUID, IPv6 Address, IPv6 Prefix, IPv6 Prefix Bits, Name, Network, Network View, Username, and User Group.
- **Object Change Network IPv4:** Disable, Network, and Network View.
- **Object Change Network IPv6:** Disable, Network, Network View, Username, and User Group.
- **Object Change Range IPv4:** Disable, Network, Network View, Server Association Type, Username, and User Group.
- **Object Change Range IPv6:** Address Type, Disable, Network, Network View, Server Association Type, Username, and User Group.
- **Object Change Host Address IPv4:** Host, IPv4 Address, MAC, Network, Network View Association Type, Username, and User Group.
- **Object Change Host Address IPv6:** Address Type, DUID, Host, IPv6 Address, IPv6 Prefix, IPv6 Prefix Bits, Network View, Username, and User Group.



- **Object Change Discovery Data:** Discoverer, IP Address, Is IPv4, Operation Type, Unmanaged, Username, and User Group.
- **Object Change DNS Record:** Auto Created Records, DNS View, Network View, Operation Type, Record Name, Record Type, User Group, Username, and Zone Name.
- **Object Change DNS Zone:** DNS View, Network View, Operation Type, User Group, Username, Zone Name, and Zone Type.

#### Note

An event may not be triggered and the template may not execute in the following scenarios:

- For events of type **Object Change DNS Record**.
- Microsoft synchronization of records, DDNS update of records, and nameserver update of records are not supported.
- For SOA records, serial numbers are not automatically incremented.
- For NS records, the **Comment** field is not updated.
- For an RPZ rule, add and update operations are not supported.
- For a shared record, add and update operations are not supported.
- For events of type **Object Change DNS Zone**, Microsoft synchronization of zones is not supported.
- Transfer of secondary zones is not supported.

4. Click **Next**. If you have selected **DNS RPZ**, or **Security ADP** or **Object Change Discovery Data** as the event type, go to Deduplicating Events below to configure deduplication. Otherwise, go to Selecting Action Template below to select an action template.

#### Note

- The event type you select here affects the templates that are available when you select the RESTful API template you want to use for the outbound notifications. For example, if you select **DNS RPZ** as the event type, only templates configured for DNS RPZ event type are available for selection.
- For the Cisco ISE endpoint, only the DNS RPZ, Security ADP, DHCP Leases, and ADP events are applicable.

## Enabling Query FQDN for Outbound Notifications

Infoblox allows you to configure support for query FQDN for outbound threat protection events and choose maximum labels in FQDN that can be configured at the Grid and/or member level. When you enable query FQDN, event data will contain the `query_fqdn` field, if any, which is limited by the domain level. The outbound template executes the parameters and fields against the notification criteria to verify if the notification rule works for the selected security ADP event type.

Note that the maximum domain level is set to three and you can query for domain levels up to three. For example, if you set the domain level to two, you can query for domain a.com, but if you query a.b.com, then the outbound template does not execute the details against the notification criteria. When you set the domain level to three, you can query for a.b.com, but if you query for a.b.c.com, then the details are not executed. Query FQDN automatically prefixes a \*. at the beginning of the domain name if the FQDN is longer.

You can enable query FQDN through the *Grid Security Properties* or *Member Security Properties* editor. A warning message is displayed if the notification rule uses Query FQDN for filtering or deduplication when it is not enabled on each member.

To enable query FQDN for outbound notifications, complete the following steps:

1. From the **Data Management** tab, select the **Security** tab, then click **Grid Security Properties** from the Toolbar.  
or  
From the **Data Management** tab, select the **Security** tab -> **Members** tab -> *member* checkbox, and then click the Edit icon.



2. In the *Grid Security Properties* or *Member Security Properties* editor, click **Toggle Advanced Mode**, select the **Ecosystem** tab, and complete the following:
  - **Enable Query FQDN for Threat Protection Events:** Select this checkbox to enable NIOS to use DNS query FQDNs for Outbound threat protection events in the Grid.
  - **Max domain level:** Select a value from the drop-down list to set the maximum domain level for query FQDNs. You can choose 2 or 3.

## Deduplicating Events

### Note

This step appears only if you have selected **DNS RPZ**, or **Security ADP**, **Object Change Discovery Data**, or **DXL Events** as the event type.

Depending on your configuration, the appliance sends the first RPZ, or threat protection, or object change discovery data event and deduplicates subsequent events that match your filtering criteria within the specified lookback interval. The hits are considered based on the following fields for each of these event types:

- **RPZ events:** Source IP, Query Name, RPZ Policy, and other related fields.
- **ADP hits:** Source IP, Rule ID, and other corresponding fields.
- **Object Change Discovery Data:** Discoverer, IP Address, and other fields.

1. To avoid excessive notifications received at the endpoint, complete the following to configure event deduplication:
  - **Enable event deduplication:** Select this to enable event deduplication for RPZ, or ADP, or data discovery hits. When you enable deduplication, the appliance suppresses redundant notifications based on your configuration.
  - **Log all dropped events due to deduplication to the syslog:** Select this if you want to log all the events that have been dropped due to deduplication. Selecting this allows the appliance to record all the dropped events to the syslog.
  - **Select the fields to use for deduplication:** From the **Available** table, pick the fields you want to use for filtering the deduplication and move them to the **Selected** table using the right arrow. You can also deselect any fields by selecting and moving them from the **Selected** table to the **Available** table using the left arrow. Event deduplication is done based on the conditions of the selected fields. The following example explains how deduplication works if two RPZ hits occur within the lookback interval, as follows:

```
RPZ hit 1 / ADP hit 1 / Data Discovery 1 : source_ip: 1.2.3.4, query_name:
server1.bad.com, rpz_policy: NXDOMAIN, query_type: qname,
network.network_view: internal, network.network: 1.2.3.0/24
```

```
RPZ hit 2 / ADP hit 2 / Data Discovery 2 : source_ip: 1.2.3.4,
query_name: www.something.com, rpz_policy: NXDOMAIN, query_type: qname,
network.network_view: internal, network.network: 1.2.3.0/24
```

If you have selected only **Source IP** for deduplication, the appliance sends only the first RPZ event to the endpoint. If you have selected both **Source IP** and **Query Name**, both RPZ events are sent to the endpoint.

- **Lookback Interval:** Enter the time interval during which the appliance evaluates RPZ hit, or ADP hit, or data discovery events and stops sending redundant events to the endpoint (based on your configuration). At the end of this interval, the appliance resumes scanning of the client IP, query FQDN, or network for RPZ events. The minimum interval is five seconds and the maximum is 15 minutes. The default is 10 minutes.
2. Click **Next** to select an action template for the endpoint, as described in the following section, *Selecting Action Template*.

## Selecting Action Template

1. In this step, select the outbound template you want to use for outbound notifications. The appliance validates the event type that is added to the notification rule and then matches that with the event type configured in the template.

In the **Template** field, click **Select Template** to associate an action template with the notification rule. If there are multiple templates, the <DXL or RESTful API> *Template Selector* dialog box is displayed, from which you can select an action template. Note that only templates that have the same event type configured for the notification rule appear in this dialog.

The following information is displayed about the selected action template:

- **Vendor Type:** The vendor type associated with the endpoint.
  - **Template Type:** The type of action that will be taken for the matching events.
  - **Parameters:** Displays the associated parameters of the template, such as **Name**, **Value**, and **Type**. You can click the **Value** cell and modify the value for the parameter.
2. Save the endpoint configuration.

## Modifying Notification Rules

To modify a notification rule, complete the following steps:


1. From the **Grid/System** tab, select the **Ecosystem** tab -> **Notification** tab, click the **Action** icon next to the notification rule and select **Edit** from the menu.
2. The *Notification Rule* editor provides the following tabs from which you can modify data:
  - **General:** You can modify the **Target** and **Comment** fields.
  - **Rules:** You can edit the event type and the rule, as described in this section, Configuring Notification Rules.
  - **Templates:** You can select a new action template for the notification rule.
3. Save the configuration.

## Viewing All Notification Rules

To view the list of notification rules, complete the following steps:

1. From the **Grid/System** tab, select the **Ecosystem** tab, and click the **Notification** tab.
2. Grid Manager displays the following information:
  - **Name:** Name of the notification rule.
  - **Target:** The target name.
  - **Action:** The action type.
  - **Comment:** Comments that were entered for the notification rule.
  - **Disable:** Displays whether the notification rule is disabled.

You can perform the following on this tab:

- Click the Action icon  and perform the following:
  - **Edit:** Select this to modify the notification rule.
  - **Delete:** Select this to delete a notification rule.
  - **Test Rule:** Select this to execute the parameters and fields of a template against the notification criteria and verify whether the notification rule works for the event (specified in the template). Make changes to the template if required, and you can view this information in the debug log. The test rules go through the following stages: filtering, enrichment, and deduplication.
  - **View Debug Log:** Select this to view debugging messages for the selected notification rule.
- Edit the notification rule information.
  - Select the notification rule, and then click the Edit icon.
- Delete a notification rule.
  - Select the notification rule, and then click the Delete icon.

Note when you remove all the notification rules associated with an endpoint, all the debug logs for that endpoint will also be removed.

- Print the list of notification rules.

- Click the Print icon.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.
- Create a quick filter to save frequently used filter criteria:
  - In the filter section, click **Show Filter** and define filter criteria for the quick filter.
  - Click **Save** and complete the configuration in the *Save Quick Filter* dialog box. The appliance adds the quick filter to the quick filter drop-down list in the panel. Note that global filters are prefixed with [G], local filters with [L], and system filters with [S].
- Sort the notification rules in ascending or descending order by column.

## Sample Templates

This section contains sample templates that illustrate the following:

- Session management using login and logout templates
- Viewing event types
- XML parsing and handling
- Steps execution through backtracking and loops in a template
- Operations using namespace variables
- IP-related XC operations
- WAPI integration
- Endpoint reference
- Active template functions

For more information about configuring outbound notification rules, templates, and endpoints, see [Outbound Notification Overview](#).



### Note

Some of the templates are specific for Rapid7 servers.

## Session Management Template

Following is a sample session management template for REST API endpoint:

```
{
  "name": "session",
  "version": "2.0",
  "type": "REST_ENDPOINT",
  "vendor_identifier": "Rapid7",
  "path": "/api/1.1/xml",
  "keepalive": true,
  "login_template": "login",
  "logout_template": "logout",
```

```
"session_duration": 60000
}
```

**Note**

The `login_template` and `logout_template` lines specify the templates to be executed for login and logout. The `session_duration` line specifies the length of the session. You must upload the login and logout templates before configuring a session management template that refers them.

Following is a sample session management template for DXL endpoint:

```
{
  "name": "session_tmplate",
  "version": "3.0",
  "type": "DXL_ENDPOINT",
  "vendor_identifier": "DXL",
  "dxl_keep_alive_interval": 60,
  "dxl_topic": "/outbound/session"
}
```

Following is a sample session management template for syslog endpoint:

```
{
  "version": "5.0",
  "vendor_identifier": "SyslogEP",
  "name": "Syslog Session",
  "type": "SYSLOG_ENDPOINT",
  "comment": "Syslog session template",
  "path": "/wapi/v2.3/",
  "override_path": true,
  "timeout": 123,
  "keepalive": true,
  "retry": 1,
  "retry_template": 0,
  "rate_limit": 200
}
```

Login Template

```
{
  "name": "login",
  "version": "2.0",
}
```

```

"vendor_identifier": "Rapid7",
"type": "REST_EVENT",
"event_type": ["SESSION"],
"content_type": "text/xml",
"quoting": "XMLA",
"steps": [
  {
    "body": "${XC:ASSIGN:{H:Authorization}:{S:}}",
    "operation": "NOP",
    "name": "login: clear basic auth"
  },
  {
    "parse": "XMLA",
    "operation": "POST",
    "no_connection_debug": false,
    "name": "login: request",
    "body_list": [
      "<?xml version=\"1.0\" encoding=\"UTF-8\"?>",
      "<LoginRequest user-id=\"${UT::USERNAME}\" password=\"${
UT::PASSWORD}\"/>"
    ]
  },
  {
    "operation": "CONDITION",
    "name": "login: errorcheck",
    "condition": {
      "statements": [
        {

```

```

        "op": "!=",
        "right": "${P:A:PARSE[[name]]}",
        "left": "LoginResponse"
    },
    {
        "op": "!=",
        "right": "1",
        "left": "${P:A:PARSE{{success}}}"
    }
],
"condition_type": "OR",
"else_eval": "${XC:COPY:{S:SESSID}:{P:PARSE{{session-id}}}",
"error": true
}
}
]
}

```



**Note**

The `else_eval` line copies the session-id from the parsed reply (if successful); otherwise, it returns an error.

Logout Template

```

{
    "name": "logout",
    "version": "2.0",
    "type": "REST_EVENT",
    "vendor_identifiler": "Rapid7",
    "event_type": ["SESSION"],

```

```

"content_type": "text/xml",
"quoting": "XMLA",
"steps": [
  {
    "parse": "XMLA",
    "operation": "POST",
    "no_connection_debug": false,
    "name": "logout: request",
    "body_list": [
      "<?xml version=\"1.0\" encoding=\"UTF-8\"?>",
      "<LogoutRequest session-id=\"${S::SESSID}\"/>"
    ]
  },
  {
    "operation": "CONDITION",
    "name": "logout: errorcheck",
    "condition": {
      "statements": [
        {
          "op": "!=",
          "right": "${P:A:PARSE[[name]]}",
          "left": "LogoutResponse"
        },
        {
          "op": "=",
          "right": "1",
          "left": "${P:A:PARSE{{success}}}"
        }
      ]
    }
  }
]

```

```

    }
  ],
  "condition_type": "OR",
  "error": true
}
}
]
}

```

## Action Template

### Sample Action Template for DXL Endpoint

```

{
  "vendor_identifier": "DXL",
  "version": "5.0",
  "name": "action",
  "type": "DXL_EVENT",
  "event_type": [
    "FIXED_ADDRESS_IPV4"
  ],
  "dxl_topic": "/outbound/action",
  "steps": [
    {
      "operation": "DXL_SEND_EVENT",
      "name": "send_ip",
      "body": "address ${E:A:values{ipv4addr}}",
      "dxl_topic": "/outbound/step"
    },
    {
      "operation": "DXL_SEND_REQUEST",
      "name": "dxl_request",
      "body_list": "address ${E:A:values{ipv4addr}}",
      "dxl_topic": "/sample/service"
    }
  ]
}

```



```
}  
]  
}
```

#### Sample Action Template for DXL Endpoint Outbound

```
{  
  "version": "5.0",  
  "name": "DXL_action_template",  
  "type": "DXL_EVENT",  
  "event_type": ["RPZ","DXL"],  
  "action_type": "RPZ Action",  
  "comment": "Outbound API phase 5",  
  "content_type": "application/json",  
  "vendor_identifier": "McAfee",  
  "steps":  
  [  
    {  
      "operation": "DXL_SEND_EVENT",  
      "name": "dxl_event",  
      "body_list": [  
        "Hello Event"  
      ],  
      "dxl_topic": "/outbound/demo"  
    },  
    {  
      "operation": "DXL_SEND_REQUEST",  
      "name": "dxl_request",  
      "body_list": [  
        "Hello Service"  
      ],  
      "dxl_topic": "/isecg/sample/service"  
    },  
    {  
      "name": "log4",  
      "operation": "NOP",  
      "body": "${XC:DEBUG:{R:}}"  
    },  
  ],  
}
```

```

{
  "name": "copy_rbody",
  "operation": "NOP",
  "body_list": [
    "${XC:COPY:{L:data}:{R:BODY}}"
  ],
},
{
  "name": "debugL",
  "operation": "NOP",
  "body": "${XC:DEBUG:{L:}}"
},
{
  "operation": "DXL_SEND_EVENT",
  "name": "resend_to_dxl",
  "body_list": [
    "Modified ${L::data}"
  ],
  "dxl_topic": "/outbound/demo"
}
]
}

```

#### Sample Action Template for REST API Endpoint

Following is a sample action template for REST API endpoint:

```

{
  "name": "action",
  "vendor_identififer": "Rapid7",
  "version": "2.0",
  "content_type": "text/xml",
  "action_type": "add network or remove IP",
  "quoting": "XMLA",
  "type": "REST_EVENT",
  "event_type": [ "FIXED_ADDRESS_IPV4", "FIXED_ADDRESS_IPV6" ],
}

```

```

"steps": [
  {
    "name": "check operation type",
    "operation": "CONDITION",
    "condition": {
      "statements": [
        {
          "op": "!=",
          "right": "${E:A:operation_type}",
          "left": "INSERT"
        },
        {
          "op": "!=",
          "right": "${E:A:operation_type}",
          "left": "DELETE"
        }
      ],
      "condition_type": "AND",
      "stop": true
    }
  },
  {
    "name": "send SiteListingRequest",
    "operation": "POST",
    "body_list": [
      "<?xml version=\"1.0\" encoding=\"UTF-8\"?>",
      "<SiteListingRequest session-id=\"${S::SESSID}\" />"
    ]
  }
]

```

```

    ],
    "parse": "XMLA"
  },
  {
    "operation": "CONDITION",
    "name": "send SiteListingRequest (error check)",
    "condition": {
      "statements": [
        {
          "op": "!=",
          "right": "${P:A:PARSE[[name]]}",
          "left": "SiteListingResponse"
        },
        {
          "op": "=",
          "right": "1",
          "left": "${P:A:PARSE{{success}}}"
        }
      ],
      "condition_type": "OR",
      "error": true,
      "else_eval": "${XC:COPY:{L:site_list}:{P:PARSE}}"
    }
  },
  {
    "operation": "CONDITION",
    "name": "check whether site list is empty",

```

```

    "condition": {
      "statements": [{
        "op": "=",
        "right": "${L:L:site_list}",
        "left": "0"
      }],
      "condition_type": "AND",
      "stop": true
    }
  },
  {
    "operation": "VARIABLEOP",
    "name": "get the next site",
    "variable_ops": [{
      "operation": "POP",
      "type": "COMPOSITE",
      "destination": "L:a_site",
      "source": "L:site_list"
    }]
  },
  {
    "operation": "CONDITION",
    "name": "check site name",
    "condition": {
      "statements": [{
        "op": "!=",
        "right": "${L:A:a_site{{name}}}",

```

```

    "left": "${E:A:values{extattrs}{r7_site}{value}}",
  ]],
  "condition_type": "AND",
  "next": "check whether site list is empty",
  "else_eval": "${XC:COPY:{L:site_id}:{L:a_site{id}}}"
}
},
{
  "parse": "XMLA",
  "operation": "POST",
  "name": "send SiteConfigRequest",
  "body_list": [
    "<?xml version=\"1.0\" encoding=\"UTF-8\"?>",
    "<SiteConfigRequest session-id=\"${S::SESSID}\" site-id=\"${L:A:site_id}\"/
  >"
  ]
},
{
  "operation": "CONDITION",
  "name": "send SiteConfigRequest (error check)",
  "condition": {
    "statements": [
      {
        "op": "!=",
        "right": "${P:A:PARSE[[name]]}",
        "left": "SiteConfigResponse"
      },
    ],
  }
}

```

```

      "op": "!=",
      "right": "1",
      "left": "${P:A:PARSE{{success}}}"
    }
  ],
  "condition_type": "OR",
  "else_eval": "${XC:COPY:{L:Site}:{P:PARSE{SiteConfigResponse}}}",
  "error": true
}
},
{
  "operation": "CONDITION",
  "name": "check operation type again",
  "condition": {
    "statements": [{
      "op": "=",
      "right": "${E:A:operation_type}",
      "left": "INSERT"
    }
  ],
  "condition_type": "AND",
  "eval":
    "${XC:COPY:{L:network}:{E:values{network}}}${XC:NETWORKTORANGE:{L:network}:
    {L:range
    }}",
  "next": "insert network"
}
},
{

```

```

    "operation": "CONDITION",
    "name": "remove ip",
    "condition": {
      "statements": [{
        "op": "=",
        "right": "${E:A:event_type}",
        "left": "FIXED_ADDRESS_IPV4"
      }],
      "condition_type": "AND",
      "eval":
        "${XC:COPY:{L:ip}:{E:values{ipv4addr}}}${XC:REMOVEIP:{L:ip}:
        {L:Site{Hosts}}}",
      "else_eval":
        "${XC:COPY:{L:ip}:{E:values{ipv6addr}}}${XC:REMOVEIP:{L:ip}:
        {L:Site{Hosts}}}"
    }
  },
  {
    "operation": "CONDITION",
    "name": "jump to send",
    "condition": {
      "statements": [{
        "op": "=",
        "right": "",
        "left": ""
      }],
      "condition_type": "AND",
      "next": "send SiteSaveRequest"
    }
  }
}

```



```

    }
  },
  {
    "operation": "VARIABLEOP",
    "name": "insert network",
    "variable_ops": [{
      "operation": "PUSH",
      "type": "COMPOSITE",
      "source": "L:range",
      "destination": "L:Site{Site}{Hosts}"
    }]
  },
  {
    "parse": "XMLA",
    "operation": "POST",
    "name": "send SiteSaveRequest",
    "body_list": [
      "<?xml version=\"1.0\" encoding=\"UTF-8\"?>",
      "<SiteSaveRequest session-id=\"${S::SESSID}\">",
      "${L:x:Site}",
      "</SiteSaveRequest>"
    ]
  },
  {
    "operation": "CONDITION",
    "name": "send SiteSaveRequest (error check)",
    "condition": {

```

```

"statements": [
  {
    "op": "!=",
    "right": "${P:A:PARSE[[name]]}",
    "left": "SiteSaveResponse"
  },
  {
    "op": "!=",
    "right": "1",
    "left": "${P:A:PARSE{{success}}}"
  }
],
"condition_type": "OR",
"error": true
}
},
{
  "operation": "CONDITION",
  "name": "check operation type once more",
  "condition": {
    "statements": [{
      "op": "!=",
      "right": "${E:A:operation_type}",
      "left": "INSERT"
    }
  ],
  "condition_type": "AND",
  "stop": true
}

```

```

    }
  },
  {
    "operation": "PUT",
    "name": "add an attribute with WAPI",
    "transport": {"path": "${E:A:values[_ref]}"},
    "wapi": "v2.6",
    "body": "{\"extattrs+\": {\"r7_added\": {\"value\": \"Added to the
Rapid7 ${UT:A:TIME}\"}}}"
  }
]
}

```

#### Sample Action Template for Syslog Endpoint

```

{
  "version": "5.0",
  "name": "syslog_action_dns_record",
  "type": "SYSLOG_EVENT",
  "event_type": ["RPZ", "DXL", "DNS_RECORD", "DNS_ZONE"],
  "action_type": "Some Action",
  "comment": "Syslog Events",
  "content_type": "application/json",
  "vendor_identifier": "syslog",
  "steps":
  [
    {
      "name": "syslog_send",
      "operation": "SYSLOG_SEND_EVENT",
      "body": "${E::object_type} ${E::operation_type} with name ${E::values{name}}
on to zone ${E::values{zone}} for member ${E::member_ip} at timestamp $
${E::timestamp}"
    }
  ]
}

```

```
]
}
```

#### Action Template: Check Operation Type

```
{
  "name": "check operation type",
  "operation": "CONDITION",
  "condition": {
    "statements": [
      {
        "op": "!=",
        "right": "${E:A:operation_type}",
        "left": "INSERT"
      },
      {
        "op": "!=",
        "right": "${E:A:operation_type}",
        "left": "DELETE"
      }
    ],
    "condition_type": "AND",
    "stop": true
  }
}
```



#### Note

The "check operation type" step checks the operation type. If it is neither INSERT nor DELETE, the template execution stops.

Action Template: Get the List of Sites

```
{
  "name": "send SiteListingRequest",
  "operation": "POST",
  "body_list": [
    "<?xml version=\"1.0\" encoding=\"UTF-8\"?>",
    "<SiteListingRequest session-id=\"${S::SESSID}\" />"
  ],
  "parse": "XMLA"
},
{
  "operation": "CONDITION",
  "name": "send SiteListingRequest (error check)",
  "condition": {
    "statements": [
      {
        "op": "!=",
        "right": "${P:A:PARSE[[name]]}",
        "left": "SiteListingResponse"
      },
      {
        "op": "=",
        "right": "1",
        "left": "${P:A:PARSE{{success}}}"
      }
    ],
    "condition_type": "OR",
  }
}
```

```

"error": true,
"else_eval": "${XC:COPY:{L:site_list}:{P:PARSE}}"
}
}

```

 **Note**

The "send SiteListingRequest" and "send SiteListingRequest (error check)" steps request the list of sites on a Rapid7 server. If the response is successful, it copies the list to the L:site\_list variable.

Action Template: Locate the Site ID by Name

```

{
  "operation": "CONDITION",
  "name": "check whether site list is empty",
  "condition": {
    "statements": [{
      "op": "==",
      "right": "${L:L:site_list}",
      "left": "0"
    }],
    "condition_type": "AND",
    "stop": true
  }
},
{
  "operation": "VARIABLEOP",
  "name": "get the next site",
  "variable_ops": [{
    "operation": "POP",

```

```

"type": "COMPOSITE",
"destination": "L:a_site",
"source": "L:site_list"
  }]
},
{
  "operation": "CONDITION",
  "name": "check site name",
  "condition": {
    "statements": [{
      "op": "!=",
      "right": "${L:A:a_site{{name}}}",
      "left": "${E:A:values{extattrs}{r7_site}{value}}"
    }],
    "condition_type": "AND",
    "next": "check whether site list is empty",
    "else_eval": "${XC:COPY:{L:site_id}:{L:a_site{{id}}}"
  }
}

```

#### NOTES:

- The "check whether site list is empty", "get the next site", and "check site name" steps form a loop for finding the site with specific name in the list of sites.
- The "check whether site list is empty" step checks to see if the list is empty. If it is empty, the site with the specific name is not found and the execution of the template stops.
- In the "get the next site" step, one element is copied to the L:a\_site variable.
- In the "check site name" step, the "name" attribute of a site ( \${L:A:a\_site{{name}}} ) is compared to the value of the "r7\_site" extensible attribute of the fixed address in ( \${E:A:values{extattrs}{r7\_site}{value}} ) . If they are the same, the site ID is stored to the L:site\_id variable ( \${XC:COPY:{L:site\_id}:{L:a\_site{{id}}}} ) . If they are not the same, the execution is continued by the "check whether site list is empty" step.

### Action Template: Get Site Configuration

```
{
  "parse": "XMLA",
  "operation": "POST",
  "name": "send SiteConfigRequest",
  "body_list": [
    "<?xml version='1.0' encoding='UTF-8'?>",
    "<SiteConfigRequest session-id='${S::SESSID}' site-id='${L:A:site_id}'>"
  ],
},
{
  "operation": "CONDITION",
  "name": "send SiteConfigRequest (error check)",
  "condition": {
    "statements": [
      {
        "op": "!=",
        "right": "${P:A:PARSE[[name]]}",
        "left": "SiteConfigResponse"
      },
      {
        "op": "!=",
        "right": "1",
        "left": "${P:A:PARSE{{success}}}"
      }
    ]
  },
},
```



```

"condition_type": "OR",
"else_eval": "${XC:COPY:{L:Site}:{P:PARSE{SiteConfigResponse}}}",
"error": true
}
}

```



#### Note

Once `site-id` is known, steps "send SiteConfigRequest" and "send SiteConfigRequest (error check)" request and store the site configuration.

#### Action Template: Distinguish Between INSERT and DELETE

```

{
  "operation": "CONDITION",
  "name": "check operation type again",
  "condition": {
    "statements": [{
      "op": "==",
      "right": "${E:A:operation_type}",
      "left": "INSERT"
    }],
    "condition_type": "AND",
    "eval":
      "${XC:COPY:{L:network}:{E:values{network}}}${XC:NETWORKTORANGE:{L:network}:
      {L:range
      }}",
    "next": "insert network"
  }
}
}

```

**Note**

The "check operation type again" step determines the operation type. If the operation is "INSERT", the network of the inserted fixed address is copied to the L:network variable ( `${XC:COPY:{L:network}:{E:values{network}}}` ) and transformed to a Rapid7 range to the L:range variable ( `${XC:NETWORKTORANGE:{L:network}:{L:range}}` ). After the range is stored, the template execution jumps to the "insert network" step.

Action Template: Delete an IP Address

```
{
  "operation": "CONDITION",
  "name": "remove ip",
  "condition": {
    "statements": [{
      "op": "=",
      "right": "${E:A:event_type}",
      "left": "FIXED_ADDRESS_IPV4"
    }],
    "condition_type": "AND",
    "eval":
      "${XC:COPY:{L:ip}:{E:values{ipv4addr}}}${XC:REMOVEIP:{L:ip}:
      {L:Site{Hosts}}}",
    "else_eval":
      "${XC:COPY:{L:ip}:{E:values{ipv6addr}}}${XC:REMOVEIP:{L:ip}:
      {L:Site{Hosts}}}"
  }
},
{
  "operation": "CONDITION",
  "name": "jump to send",
```

```

    "condition": {
      "statements": [{
        "op": "=",
        "right": "",
        "left": ""
      }],
      "condition_type": "AND",
      "next": "send SiteSaveRequest"
    }
  }
}

```

#### NOTES:

- If the operation is not "INSERT" (i.e. "DELETE"), the "remove ip" step is executed. The step determines the type of fixed address. The corresponding address ( {E:values{ipv4addr}} or {E:values{ipv6addr}} ) is copied to the L:ip variable, and then the L:ip address is removed from the list of hosts in the site ( \${XC:REMOVEIP:{L:ip}:{L:Site{Hosts}}} ).
- The step "jump to send" skips the "inserting" step and jumps directly to the "send SiteSaveRequest" step.

#### Action Template: Add an IP Range

```

{
  "operation": "VARIABLEOP",
  "name": "insert network",
  "variable_ops": [{
    "operation": "PUSH",
    "type": "COMPOSITE",
    "source": "L:range",
    "destination": "L:Site{Site}{Hosts}"
  }]
}

```

The step "insert network" pushes the L:range to the list of site's "Hosts".

### Action Template: Saving New Configuration

```
{
  "parse": "XMLA",
  "operation": "POST",
  "name": "send SiteSaveRequest",
  "body_list": [
    "<?xml version='1.0' encoding='UTF-8'?>",
    "<SiteSaveRequest session-id='${S::SESSID}'>",
    "${L:x:Site}",
    "</SiteSaveRequest>"
  ],
  },
  {
    "operation": "CONDITION",
    "name": "send SiteSaveRequest (error check)",
    "condition": {
      "statements": [
        {
          "op": "!=",
          "right": "${P:A:PARSE[[name]]}",
          "left": "SiteSaveResponse"
        },
        {
          "op": "=",
          "right": "1",
          "left": "${P:A:PARSE{{success}}}"
        }
      ]
    }
  }
}
```

```

],
"condition_type": "OR",
"error": true
}
}

```



**Note**

The "send SiteSaveRequest" and "send SiteSaveRequest (error check)" steps save new site configuration to the Rapid7 server.

Action Template: Add Extensible Attributes Using WAPI

```

{
  "operation": "CONDITION",
  "name": "check operation type once more",
  "condition": {
    "statements": [{
      "op": "!=",
      "right": "${E:A:operation_type}",
      "left": "INSERT"
    }],
    "condition_type": "AND",
    "stop": true
  }
},
{
  "operation": "PUT",
  "name": "add an attribute with WAPI",
  "transport": {"path": "${E:A:values[_ref]}"},

```

```

    "wapi": "v2.6",
    "body": "{\"extattrs+\": {\"r7_added\": {\"value\": \"Added to the
Rapid7 ${UT:A:TIME}\"}}}"
  }

```

**NOTES:**

- The "check operation type once more" step determines the operation type. If the operation is not "INSERT" (i.e. "DELETE"), the template execution stops.
- For the "INSERT" operation, the last step "add an attribute with WAPI" is executed. This step adds an extensible attribute "r7\_added" to the fixed address through RESTful API. The value of the attribute has the current timestamp ( \${UT:A:TIME} ).

Action Template: Add a Host Record

The following sample template, if assigned to a DHCP network notification rule, will insert a host record for any added network that matches the rule, with a hostname and domain name set by extensible attributes in the network. Detailed explanations about this sample are included in Action Template with Comments: Add a Host Record below.

```

{
  "version": "1.0",
  "name": "Insert host record",
  "comment": "Will automatically insert a host record for new network
insertions, assumes the network has a 'Zone' extensible attribute,
optionally a 'Hostname' extensible attribute as well",
  "type": "REST_EVENT",
  "event_type": [
    "NETWORK_IPV4"
  ],
  "action_type": "Insert a host record",
  "vendor_identifier": "WAPI 2.3",
  "transport": {
    "content_type": "application/json",
  },
  "steps":

```

```

[
  {
    "name": "stop if it is not a network insert",
    "operation": "CONDITION",
    "condition": {
      "condition_type": "AND",
      "statements": [
        {
          "left": "${E:A:operation_type}",
          "op": "!=",
          "right": "INSERT"
        }
      ],
      "stop": true
    }
  },
  {
    "name": "stop if we don't have the zone EA set, else save it",
    "operation": "CONDITION",
    "condition": {
      "condition_type": "AND",
      "statements": [
        {
          "left": "${E:A:values{extattrs}{Zone}{value}}",
          "op": "==",
          "right": ""
        }
      ]
    }
  }
]

```

```

    ],
    "stop": true,
    "else_eval": "${XC:COPY:{L:ZONE}:{E:values{extattrs}{Zone}
{value}}}"
  }
},
{
  "name": "get the hostname or use a default value",
  "operation": "CONDITION",
  "condition": {
    "condition_type": "AND",
    "statements": [
      {
        "left": "${E:A:values{extattrs}{Hostname}{value}}",
        "op": "!=",
        "right": ""
      }
    ],
    "eval":
      "${XC:COPY:{L:HOSTNAME}:{E:values{extattrs}{Hostname}{value}}}",
      "else_eval": "${XC:ASSIGN:{L:HOSTNAME}:{S:defaulthostname}}"
  }
},
{
  "name": "insert the host record with the next available IP",
  "operation": "POST",
  "transport": {
    "path": "record:host"
  }
}

```



```

    },
    "body_list": [
        "{",
        "\"name\": \"${L:A:HOSTNAME}.${L:A:ZONE}\",",
        "\"ipv4addrs\": [{\"ipv4addr\":",
        "\"func:nextavailableip:${E:A:values{network}}\"}],",
        "\"comment\": \"Inserted via outbound\"",
        "}"
    ]
}
]
}

```

#### Action Template with Comments: Add a Host Record

The comments (labeled as **Note**;) embedded in the sample template explain the operation for each section of the template. Note that the execution of this sample template will cause a single **POST** request to be sent with the following body for an insertion of a **10.0.0/24** network with the Zone extensible attribute set to test.com and the Hostname extensible attribute set to **name** :

```

{"name": "name.test.com","ipv4addrs": [{"ipv4addr":
"func:nextavailableip:10.0.0.0/24"}],"comment": "Inserted via a template"}

```

---

**Note:** The preamble of the template specifies the version, version, name and other relevant fields.

---

```

{
  "version": "1.0",
  "name": "Insert host record",
  "comment": "Will automatically insert a host record for new network
insertions,assumes the network has a 'Zone' extensible attribute, optionally
a 'Hostname' extensible attribute as well",
  "type": "REST_EVENT",

```

---

**Note:** The `event_type` field specifies that this template is used for IPv4 network events only.

---

```
"event_type": [  
  "NETWORK_IPV4"  
],
```

---

**Note:** The `action_type` and `vendor_identifier` fields describe the template type and the vendor type.

---

```
"action_type": "Insert a host record",  
"vendor_identifier": "WAPI 2.3",
```

---

**Note:** The following specifies that the template is going to send JSON to the server.

---

```
"transport": {  
  "content_type": "application/json",  
},
```

---

**Note:** The following are steps that will be executed sequentially.

---

```
"steps":  
[
```

---

**Note:** The first step will stop the execution, without an error, if the network event received is not an insertion – it could be a modify, for example.

---

```
{
  "name": "stop if it is not a network insert",
  "operation": "CONDITION",
  "condition": {
```

---

**Note:** You can specify only one statement: either AND or OR would work.

---

```
    "condition_type": "AND",
    "statements": [
```

---

**Note:** The match is to check if `operation_type` in the event is different from INSERT. It is a good practice to put the event variable on the left side so if it is not present in the event, the template would not fail.

---

```
      {
        "left": "${E:A:operation_type}",
        "op": "!=",
        "right": "INSERT"
      }
    ],
```

---

**Note:** The directive that stops the execution if the condition matches.

---

```
      "stop": true
    }
  },
```

---

**Note:** The second step will stop the execution if the inserted network does not have the Zone extensible attribute configured. If it has the extensible attribute, it will be put in a temporary local variable for easier access later on.

---

```
{  
  "name": "stop if we don't have the zone EA set, else save it",  
  "operation": "CONDITION",  
  "condition": {
```

---

**Note:** Similar to the previous section, we have only one statement, so either AND or OR would work. The condition ensures that the Zone extensible attribute is set to a value.

---

```
    "condition_type": "AND",  
    "statements": [  
      {
```

---

**Note:** As previously mentioned, non-existent variable access in the left side of a condition will not cause an error, but instead return an empty value.

---

```
        "left": "${E:A:values{extattrs}{Zone}{value}}",  
        "op": "==",  
        "right": ""  
      }  
    ],
```

---

**Note:** If the extensible attribute is empty or nonexistent, the operation should stop here.

---

```
      "stop": true,
```

---

**Note:** Otherwise, it will copy the zone value to the local ZONE variable.

---

```
    "else_eval": "${XC:COPY:{L:ZONE}:{E:values{extattrs}{Zone}{value}}}"  
  }  
},
```

---

**Note:** This step is similar to the zone step above. However, if the host name is not set, it will put a default host name in a local variable to provide the default.

---

```
{  
  "name": "get the hostname or use a default value",  
  "operation": "CONDITION",  
  "condition": {  
    "condition_type": "AND",  
    "statements": [  
      {  
        "left": "${E:A:values{extattrs}{Hostname}{value}}",  
        "op": "!=",  
        "right": ""  
      }  
    ],  
  },  
}
```

---

**Note:** This is executed if the extensible attribute is present, by copying its value to HOSTNAME.

---

```
"eval": "${XC:COPY:{L:HOSTNAME}:{E:values{extattrs}{Hostname}{value}}}",
```

---

**Note:** Otherwise, the following is executed if the extensible attribute is empty or not present, by assigning the `defaulthostname` string instead.

---

```
"else_eval": "${XC:ASSIGN:{L:HOSTNAME}:{S:defaulthostname}}"  
}  
},
```

---

**Note:** This step will finally contact the endpoint and in this case insert the host.

---

```
{  
  "name": "insert the host record with the next available IP",
```

---

**Note:** This defines the HTTP operation to use.

---

```
"operation": "POST",
```

---

**Note:** The endpoint is configured starting with `https://master_ip/...` The endpoint template overrides the path with `/wapi/v2.3/` so by default all template requests would go to `https://master_ip/wapi/v2.3/`. In this step, we want to insert a host, so `record:host` is appended to the URI above (no override is set here) to arrive to the valid RESTful URI `https://master_ip/wapi/v2.3/record:host`.

---

```
"transport": {  
  "path": "record:host"  
},
```

---

**Note:** This is the text that will be sent to the server in the POST's BODY.

---

```
"body_list": [  
  "{",
```

---

**Note:** This references the local variables that were previously assigned.

---

```
"\"name\": \"${L:A:HOSTNAME}.${L:A:ZONE}\"",
```

---

**Note:** This option signifies that the RESTful API will use the next available IP in the network as the address for this host.

---

```
"\"ipv4addrs\": [{\"ipv4addr\":  
  \"func:nextavailableip:${E:A:values{network}}\"}],",
```

---

**Note:** This is the comment stating that the action is done through the RESTful API template.

---

```
"\"comment\": \"Inserted via outbound\"",  
  "}"  
]  
}  
]  
}
```

### Endpoint reference

`UT::ENDPOINT` contains WAPI reference to the current endpoint object that can be used by WAPI steps. You can make changes to the endpoint configuration during a template execution.

**Example:**

```
{  
  "version": "4.0",  
  ...
```

```

"steps": [
...
{
"operation": "PUT",
"name": "update_endpoint",
"transport": {"path": "${UT::ENDPOINT}"},
"wapi": "v2.7",
"body": {"comment": "outbound"}
},
...
],
...
}

```

#### Active Template Functions

##### Example:

```

{
"version": "4.0",
"name": "example",
"event_type": ["RPZ"],
"type": "REST_EVENT",
"functions": {
"functions": {
"is_ipv4_address": {
"steps": [
{
"operation": "CONDITION",
"condition": {

```



```

"statements": [
{
"left": "${L:A:address}",
"op": "!~",
"right": ":"
}
],
"condition_type": "AND",
"eval": "${XC:ASSIGN:{L:result}:{B:true}}",
"else_eval": "${XC:ASSIGN:{L:result}:{B:false}}"
}
}
]
}
},
"steps": [
...
{
"operation": "FUNCTION",
"function_name": "is_ipv4_address",
"body": "${XC:ASSIGN:{L:address}:{S:10.0.0.1}}"
},
...
]
}

```

## Configuring BloxOne Threat Defense Cloud Clients for Outbound

The Infoblox BloxOne Threat Defense Cloud Client on NIOS allows the interaction between BloxOne Threat Defense Cloud and outbound endpoints so you can collect blocked/logged request via feeds or domains detected by Threat

Insight in BloxOne Threat Defense Cloud and send the outbound events to external endpoints. When you enable and configure BloxOne Threat Defense Cloud Client on an on-prem NIOS member, the client uses threat API calls to request RPZ events from BloxOne Threat Defense Cloud, and then convert the data into outbound events. These events are periodically synchronized (between BloxOne Threat Defense Cloud and NIOS) and sent to the configured outbound endpoints. Note that the client requests only subsequent data since the last data timestamp, and each synchronization happens based on the schedule and retrieves only the current data.

You can configure notification rules to filter incoming events using the following fields: Threat Origin (NIOS, BloxOne Threat Defense Cloud), BloxOne Threat Defense Cloud Hit Type (DNS RPZ, Threat Analytics), BloxOne Threat Defense Cloud Hit Class and BloxOne Threat Defense Cloud Hit Property. When you configure notification rules to filter incoming events using these fields for BloxOne Threat Defense Cloud Client, relevant information gets synchronized with the event types that you add to the list. This synchronization happens periodically based on the interval that you define. For more information about notification rules, see [Configuring Notification Rules](#).

You can select any Grid member to execute the BloxOne Threat Defense Cloud Client. Infoblox uses event filters on the selected Grid Member to limit the amount of logs. For debugging purposes, information about the client connection status will be displayed in the infoblox.log file. An error is logged in the debug mode for any exceptions that appear when the data is requested and received from the BloxOne Threat Defense Cloud. NIOS logs any critical messages in the syslog.

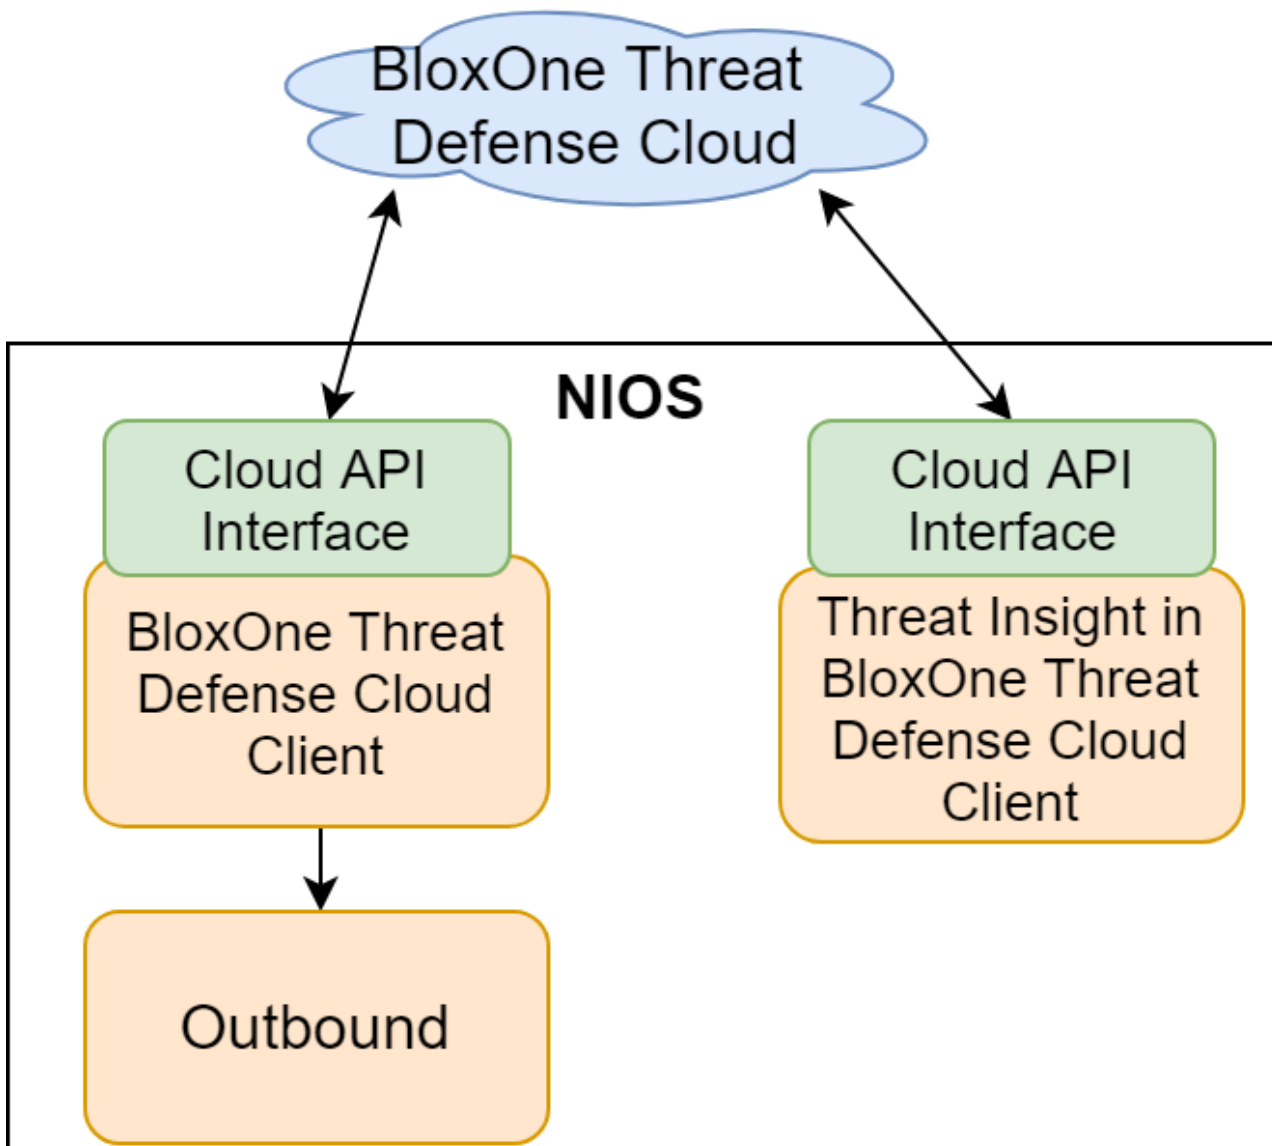
You must specify the email address and password in the *Grid Properties Editor* before you enable the BloxOne Threat Defense Cloud Client. For more information about configuring Integration with BloxOne threat defense cloud, see below. The server stores the email address and the password so that it can request a new API key. The server requests an API key through the Cloud Services Portal, so that the cloud client is authorized to retrieve data from BloxOne Threat Defense Cloud.



#### Note

Before you configure the BloxOne Threat Defense Cloud Client for outbound, ensure that you have installed the **Security Ecosystem** license.

The following figure shows how Threat Insight in the BloxOne Threat Defense Cloud client and BloxOne Threat Defense Cloud Client use a common API interface to interact with BloxOne Threat Defense Cloud. For more information about enabling BloxOne threat defense cloud client for outbound, see below.



### Best Practices for Configuring BloxOne Threat Defense Cloud Client

- Ensure that you have enabled the following on the BloxOne Threat Defense Cloud Client:
  - An email address and a password.
  - A Grid member that is online.
- Ensure that at least one outbound notification rule for DNS RPZ event type is active for outbound settings.
- Only superusers can update the BloxOne Threat Defense Cloud Client settings.
- If the timestamp for the data collected by the BloxOne Threat Defense Cloud Client is ahead of the current time in NIOS, then such events are logged in the syslog. In such an instance, the client does not request any data until the current time reaches the timestamp of the data that is collected and it logs a message in the Infoblox.log based on the time interval that you have set.

### Configuring Integration with BloxOne Threat Defense Cloud

To integrate the BloxOne Threat Defense Cloud client with BloxOne Threat Defense Cloud, you must have already created a user profile and the API key for the user profile in the Cloud Services Portal.

To configure the BloxOne Threat Defense Cloud client to integrate with BloxOne Threat Defense Cloud, you must configure the URL of the Cloud Services Portal and credentials for logging in to the portal. Complete the following steps:

1. **Grid:** From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.  
**Standalone appliance:** From the **System** tab, select the **System Manager** tab, and then select **System Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties Editor* or the *System Properties Editor*, click **Toggle Advanced Mode** to switch to the advanced mode.  
Note that if the editor is already in the advanced mode, then you will see the Toggle Basic Mode button.
3. On the **BloxOne Threat Defense Cloud Integration** tab -> **Basic** tab, specify the following in the *BloxOne Threat Defense Cloud Integration* section:  
**BloxOne Threat Defense Cloud Integration**
  - **URL:** Displays the REST API URL of the Infoblox Cloud Services Portal.
  - **Credentials:**
    - **Email:** Enter the email address that is registered in the Cloud Services Portal. This email address is used for authorization by the Cloud Services Portal.
    - **Password:** Enter the password that is registered in the Cloud Services Portal. This password is used for authorization by the Cloud Services Portal.
    - **Test Connection:** Click this to test the connectivity between NIOS and the Cloud Services Portal.
4. Save the configuration.

## Enabling BloxOne Threat Defense Cloud Client for Outbound

To configure an BloxOne Threat Defense Cloud Client to collect event types from BloxOne Threat Defense Cloud and send them to external endpoints, complete the following steps:

1. From the **Grid** tab, select the **Ecosystem** tab -> **Outbound Endpoint** tab, and then click **BloxOne Threat Defense Cloud Client** from the Toolbar.
2. In the *BloxOne Threat Defense Cloud Client* editor, complete the following:
  - **Enable Cloud Client:** Select this checkbox to enable the BloxOne Threat Defense Cloud Client to send outbound events.
  - **Grid member:** Click **Select** to select a Grid member on which you run the configured client. Click **Clear** to clear the value. You can select any Grid member where the cloud client must be executed.
  - **Interval:** Specify how often to request the list of event types from BloxOne Threat Defense Cloud, in seconds or minutes. This value is set to one minute by default. The time interval is measured from the previous data synchronization.
  - **The list of requested event types:** Select the respective checkbox to enable or disable an event type. The event types that you request from the BloxOne Threat Defense Cloud are listed here. You cannot add or remove them.
3. Save the configuration.

## Configuring Outbound Cloud Clients

The Infoblox outbound cloud client allows interaction of the Cloud Services Portal with external endpoints. You can use the outbound cloud client, which in turn uses threat API calls, to request security events from the Cloud Services Portal and convert data to outbound events. These events are periodically published to outbound framework for filtering and enrichment before they are sent to the configured external endpoints. The outbound cloud client stores the last timestamp from data that is received from the Cloud Services Portal. For the next security event request, it requests data from the last timestamp until the current time. It does not request any historical data from the cloud service portal.

With the outbound cloud client, you can request security data such as DNS RPZ, threat analytics, content/domain category, a class/property, and a client type. Infoblox enables you to configure notification rules to filter incoming events using the following fields: Threat Origin (NIOS, Cloud), Threat Source (DNS RPZ, Threat Analytics, Content Category), Threat Class and Threat Property. For more information, see [Configuring Notification Rules](#).

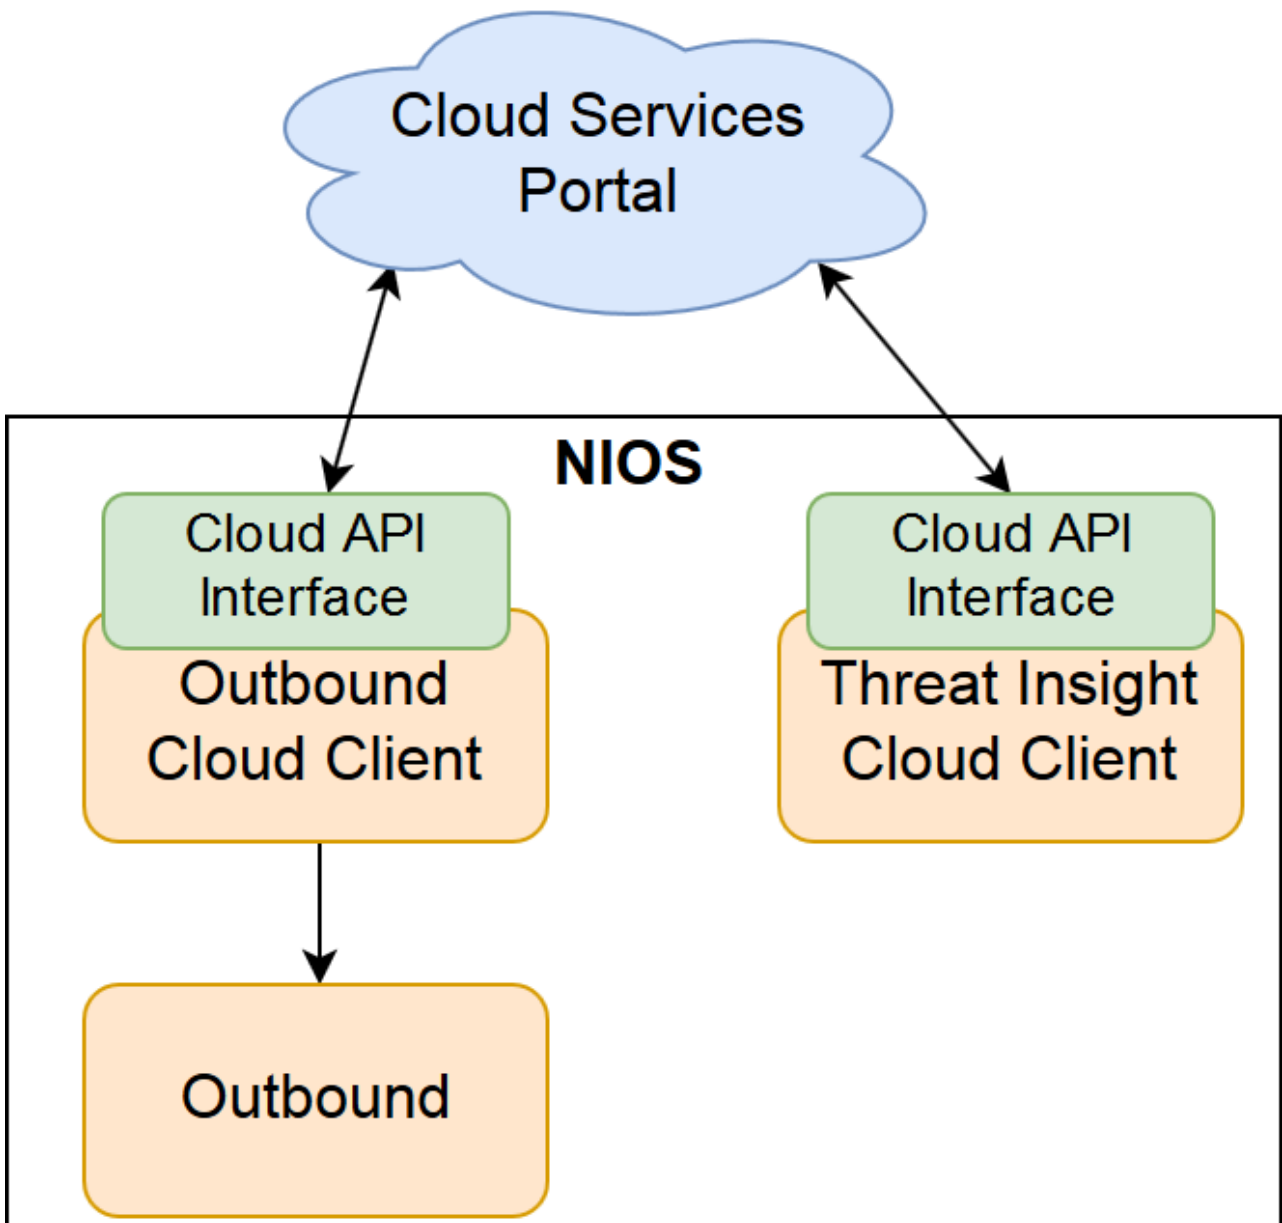
You must enable either the username or password and define an API key to enable a client. A client can be a threat insight in the cloud client or an outbound cloud client. The server stores the username and password to request a new API key when the existing API key expires. Note that all clients use a common API interface.

You can select any Grid member to execute the outbound cloud client. Infoblox uses event filters on the selected Grid Member to limit the amount of logs. The audit log displays information about the client connection status, exceptions during connection and data request, and data received from the Cloud Services Portal. NIOS logs any critical messages in the syslog.



**Note**

Before you configure the outbound cloud client, ensure that you have installed the Response Policy Zones and Security Ecosystem licenses.



## Best Practices for Configuring Outbound Cloud Clients

- Ensure that you have enabled the following on the outbound client:
  - Define either a username or a password.
  - Enable at least one event type.
  - Specify a Grid member and ensure that the Grid member is online.
- Any changes to the API request or response format affects the feature.
- At least one outbound notification rule for DNS RPZ event type must be active for outbound settings.
- You cannot insert or delete outbound cloud clients.
- Only superusers can update an object.

## Configuring Outbound Cloud Clients

To configure BloxOne Threat Defense Cloud API, complete the following:

1. From the **Grid** tab, select the **Grid Manager** tab, and then select **Grid Properties** -> **Edit** from the Toolbar.
2. In the *Grid Properties Editor*, click **Toggle Advanced Mode**, and click the **BloxOne Threat Defense Cloud Integration** tab.
3. Specify the following in the **Basic** tab:
  - BloxOne Threat Defense Cloud API**
    - **URL**: Displays the REST API URL of the Infoblox BloxOne Threat Defense Cloud Services Portal. You can use the `set cloud_services_portal` CLI command in maintenance mode to update the REST API URL.
    - **Credentials**
      - **Email**: Enter the email address that is registered in the Infoblox Cloud Services Portal.
      - **Password**: Enter the password that is registered in the Infoblox Cloud Services Portal.
    - **Test Connection**: Click this to test the connectivity between the outbound cloud client and the BloxOne Threat Defense Cloud Services Portal.
4. Save the configuration.

### 3.3.2 Threat Insight in the Cloud Client Editor (update)

User will not be able to specify API key from the Editor. It should be removed.

## Enabling Client for Outbound Events in the Cloud

To modify an outbound cloud client:

1. From the **Grid** tab, select the **Ecosystem** tab -> **Outbound Endpoint** tab, and then click **BloxOne Threat Defense Cloud Client** from the Toolbar.
2. In the *BloxOne Threat Defense Cloud Client* editor, complete the following:
  - **Enable Cloud Client**: Select this checkbox to enable the outbound events in the cloud client.
  - **Grid member**: Click **Select** to select a Grid member to execute the outbound cloud client. Click **Clear** to clear the value. You can select any Grid Member where the cloud client must be executed.
  - **Interval**: Specify how often to request outbound events detected in the cloud client in seconds or minutes. This value is set to one minute, by default. The time interval is measured from the previous request cycle.
  - **The list of requested event types**: Select the respective checkbox to enable or disable an event type. The event types that you request from the Cloud Services Portal are listed here. You cannot add or remove them.
3. Save the configuration.

## Infoblox Subscriber Services

 **Important**

Infoblox Subscriber Services is not supported in NIOS 8.6.0. Although Subscriber Services is supported in NIOS 8.6.1, Infoblox recommends that you do not use it in this version.

The Infoblox Subscriber Services provides a scalable, enterprise-grade solution that provides visibility to subscriber activities and complete filtering capabilities by combining advanced DNS services with subscriber identification, threat protection policies, and MSP (Multi-Services Proxy). The Infoblox Subscriber Services solution includes the following:

- **Infoblox Subscriber Insight** - Infoblox Subscriber Insight automates the process of identifying infected subscriber devices that are trying to connect to malicious domains. This solution augments the malware incident logs with the subscriber identity information received via RADIUS accounting messages and generates a report to display RPZ violations per subscriber ID. You can also identify subscribers who access specific domains for purposes other than security.
- **Infoblox Subscriber Policy Enforcement** - Infoblox Subscriber Policy Enforcement enables the selection of applicable policies for the subscriber. Policies are any combinations of RPZs. You can use this product to create value-added service plans or packages for different subscribers.
- **Infoblox Subscriber Parental Control** - Infoblox Subscriber Parental Control enables subscribers to manage Internet access and content for their mobility devices, houses, families, or corporations. Subscribers can restrict or allow access to content based on content categories and domains.

This section includes the following topics:

- [Infoblox Subscriber Insight and Subscriber Policy Enforcement](#)
- [Infoblox Subscriber Parental Control](#)
- [Scaling Using Subscriber Sites](#)
- [Managing AVPs \(Attribute Value Pairs\)](#)
- [Monitoring Subscriber Policy Violations](#)

## Infoblox Subscriber Insight and Subscriber Policy Enforcement

The Infoblox Subscriber Insight solution provides a mechanism to monitor events related to the subscriber session. This solution allows you to identify subscriber devices, such as laptops, computers, tablets, and smartphones on your data networks that are violating RPZ rules. It can also find the type of domain the subscriber tries to access.

The Infoblox Policy Enforcement solution analyzes the DNS queries, identify subscribers and correlates the information to enforce the subscriber security policies per subscriber.

The solution works by receiving RADIUS accounting messages from a RADIUS server through the NAS (Network Access Server) gateway. The DNS server caches the RADIUS accounting messages, which includes subscriber information and subscriber security policies. The subscriber security policy specifies the RPZs that are applicable for a subscriber. DNS RPZs are used to determine the bad FQDNs. If a subscriber, who has opted for the service, queries an FQDN that is listed in the RPZ, the DNS resolver performs RPZ actions for the subscriber query.

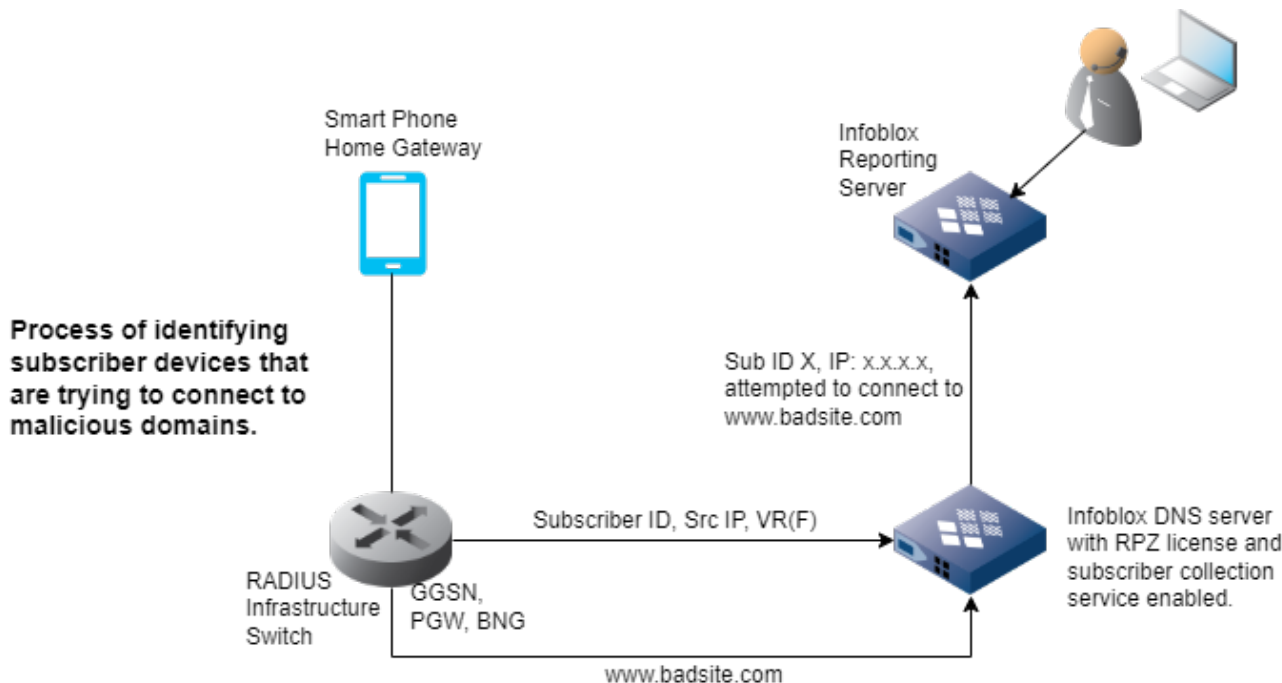
Subscribers behind a home gateway network are identified by their local ID or client ID, which is the MAC address of the subscriber device. The local ID is received by the DNS server as part of the RADIUS accounting message. As all subscribers behind a home gateway network will have the same IP address of the home router, the local ID is used to create separate records for each of the subscribers in the subscriber cache. Each subscriber behind the home gateway network can have their own policy. Note that if a guest connects through a home gateway network, then the default home router policy is applied to the guest device.

The Infoblox Subscriber Insight and Subscriber Policy Enforcement is currently supported on the following Infoblox appliances: IB-1415, IB-1425, IB-2215, IB-2225, PT-1405, PT-2205, IB-4030-10GE, IB-VM-1405, IB-VM-1415, IB-VM-1425, IB-VM-2205, IB-VM-2225, IB-VM-1425, and IB-FLEX. You can enable DNS cache acceleration feature on IB-4030-10GE, and IB-FLEX appliances.

As illustrated in the figure below, the Infoblox DNS member, with subscriber collection service enabled, receives RADIUS accounting messages, which includes subscriber information and the subscriber security policies, through the NAS gateway. The subscriber security policy specifies the RPZs that are applicable for a subscriber. When a subscriber queries an FQDN that is listed in the RPZ, the DNS resolver performs RPZ actions for the subscriber query. The NIOS appliance logs all RPZ related events, conformed to CEF (Common Event Format), in the syslog. The CEF logs include the subscriber identity information, thus identifying the subscribers that are violating RPZ rules. In a Grid with a reporting

server, you can view the *Detailed RPZ Violations by Subscriber ID* report that contains information about RPZ hits by the users. For information about detailed RPZ violations by subscriber ID, see [About Dashboards](#).

#### Infoblox Subscriber Insight and Subscriber Policy Enforcement



## License Requirements and Admin Permissions

To configure Infoblox Subscriber Services, you must install the **RPZ** license. But for IB-FLEX members, the **RPZ** license is included in the **FLEX Grid Activation** license. For information about how to install licenses, see [Managing Licenses](#).

Only superusers can configure Infoblox Subscriber Services. Limited-access admin groups can perform this operation only if their administrative permissions are defined. For information about administrative permissions, see [About Administrative Permissions](#).

## Guidelines for Using Infoblox Subscriber Insight and Subscriber Policy Enforcement

Following are some guidelines to take into consideration when using Infoblox Subscriber Threat Insight and Subscriber Policy Enforcement:

- A Grid member can be associated with only one subscriber site.
- The Grid members in the subscriber site must be added as primary or secondary name servers for all RPZs.
- The subscriber collection service does not support an IPv6 only Grid member.
- All NAS gateways in a subscriber site must be configured for IPv4 only and must use the same port.
- The subscriber data is not persistent and will be cleared from the subscriber cache if you stop the subscriber collection service on all the members of a subscriber site.
- The NAS gateways can send the RADIUS accounting messages to only one Grid member (collector member) in a subscriber site.
- Overlapping networks are not supported in the Subscriber Insight and Subscriber Policy Enforcement solutions.

## Configuring Infoblox Subscriber Insight and Subscriber Policy Enforcement

To set up Infoblox Subscriber Insight and Subscriber Policy Enforcement, you must install a Grid-wide or a member level **RPZ** license and configure the Grid members to serve recursive DNS queries. Note that for IB-FLEX members, you do



not need to install an **RPZ** license as the **FLEX Grid Activation** license includes the **RPZ** license. You must also configure a subscriber site and add Grid members (collector members and RPZ members) and NAS (Network Access Server) gateways to receive RADIUS accounting messages from a RADIUS server. The RADIUS accounting messages include subscriber information (such as subscriber source IP address, subscriber ID, local ID for networks with overlapping IP addresses) and subscriber security policies. The source IP address, subscriber ID, local ID, and the subscriber security policy is mapped in the DNS cache. The collector member caches the subscriber information and the policies, which are replicated to all Grid members within the subscriber site. The RPZ members in the subscriber site applies policies for incoming subscriber queries and performs RPZ actions.

The DNS Cache Acceleration processes incoming EDNS0 packets that contains the local ID. The NIOS appliance matches the local ID against the DNS server as part of the RADIUS accounting message and populates the subscriber cache in DNS Cache Acceleration with the parental control policy information. The DNS Cache Acceleration answers all queries that comes from the DNS Cache Acceleration for each of these subscribers listed in the subscriber cache. These changes are valid for individual IP addresses with local ID only and the subnet local ID is considered as 0. For more information, see [Using the NIOS CLI](#) and for viewing the DNS accelerator cache, [Clearing DNS Cache](#) in the *Infoblox DNS Cache Acceleration Administrator Guide*.

Note the following while configuring Infoblox Subscriber Insight and Subscriber Policy Enforcement:

- You can configure up to a total of 32 RPZs in the default view and set the priorities for the RPZs. Subscribers with subscriber security policies can set the policy to any of the 32 RPZs and the RPZs are applied to selective subscribers depending on the subscriber security policies. But for subscribers who have not opted for the service, only the first five (top priority) RPZs are applied. Note that once the first hit matches, the rest of the RPZs will not be looked up.
- For DNS queries received from unknown subscriber source IP addresses, the DNS server processes the queries based on the standard DNS query processing.

To configure the Infoblox Subscriber Insight and Subscriber Policy Enforcement on supported Infoblox appliances, complete the following:

1. Obtain and install a valid Grid-wide or member level **RPZ** license. But for IB-FLEX members, the **RPZ** license is included in the **FLEX Grid Activation** license. For information about licenses, see License Requirements and Admin Permissions below. You can also configure a reporting appliance in the Grid to see subscriber reports that contain statistics about RPZ related events.
2. Configure admin permissions so admin users can manage the Infoblox Subscriber Service related tasks. For information about how to configure admin permission, see [About Administrative Permissions](#).
3. Create a subscriber site with at least one Grid member and a NAS gateway. For information, see [Adding Subscriber Sites](#). It is recommended to add more than one Grid member to the subscriber site for redundancy. You can add a maximum of five Grid members to the subscriber site.
4. Start the subscriber collection service on all the members in the subscriber site, as described in Starting and Stopping the Subscriber Collection Service below.
5. Create RPZs in the default DNS view and specify the order of RPZs. Note that only the default DNS view is supported for configuring RPZs for Subscriber Services. For information about creating RPZs, see [Configuring Local RPZs](#). For information about specifying the order of RPZs, see [Reordering RPZs](#). You can create a total of 32 RPZs. Subscribers with subscriber security policies can set the policy to any combination of the 32 RPZs and the RPZs are applied to selective subscribers depending on the subscriber security policies. But for subscribers who have not opted for the service, only the first five (top priority) RPZs are applied. The NAS gateways must provide the subscriber security policies that enable the selection of RPZs applicable for the subscriber.
6. Enable all the members in the Grid to respond to recursive queries, as described in [Enabling Recursive Queries](#).
7. Enable RPZ logging in the *Member DNS Properties* editor for each member of the subscriber site, to ensure that all events related to RPZ are logged to the syslog. Note that you can also enable logging of queries and responses, but it might significantly affect system performance. For information about how to set logging categories, see [Using a Syslog Server](#).
8. After completing the DNS configuration on the Grid members, start the DNS service on the Grid members. For information about how to start and stop the DNS service, see [Starting and Stopping the Discovery Service](#). Ensure that you enable IPv6 on the Grid members to support IPv6 subscribers. For information, see [Configuring IPv6 on a Grid Member](#).

9. Add AVPs that are not available in the list of predefined AVPs. For information, see [Adding AVPs](#).
10. Configure the subscriber ID settings to associate an AVP with the subscriber in the Subscriber Services Properties editor, as described in Configuring Subscriber Services Properties below.

After you set up the Infoblox Subscriber Insight and Subscriber Policy Enforcement, you can do the following:

- View the subscriber sites, as described in [Viewing Subscriber Sites](#).
- View the NAS gateway message rates for the accounting servers of the subscriber site, as described in [Viewing NAS Gateway Message Rates](#).
- Monitor RPZ related events and subscriber policy violations using predefined reports and the syslog, as described in [Monitoring Subscriber Policy Violations](#).
- Configure Infoblox Subscriber Parental Control solution, as described in [Infoblox Subscriber Parental Control](#).

## Starting and Stopping the Subscriber Collection Service

To start the subscriber collection service, you must have at least one **RPZ** license installed (it can be a Grid-wide license or a member-level license.) You can also stop the service when necessary.

To start or stop the subscriber collection service:

1. From the **Grid** tab -> **Grid Manager** tab -> **Services** tab, click the **Subscriber Collection** link. Grid Manager displays only the members that are running the subscriber collection service. Select the member checkbox.
2. From the Toolbar, click **Start** to start the service or **Stop** to stop the service.



### Note

The subscriber data is not persistent and will be cleared from the subscriber cache if you stop the subscriber collection service on all the members of a subscriber site.

## Configuring Subscriber Services Properties

To configure the subscriber services properties:

1. From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab, expand the Toolbar and click **Subscriber Services Properties**.
2. In the **General** tab, complete the following:
  - **Subscriber ID**: Select the AVP that you want to use as the subscriber ID. Depending on the AVP you select in this field, the appliance extracts those AVPs from the RADIUS accounting messages. All the sites in the Grid will use the AVP selected in this field as the subscriber ID. You can select one of the following AVPs from the drop-down list: APN, Calling-Station-Id, Class, IMEI, IMSI, MSISDN, and User-Name.  
Note that for subscribers who have opted for Parental Control service, you can select Calling-Station-Id as the subscriber ID for mobile networks and for fixed line, you can select User-Name as the subscriber ID.
  - **REGEX for Subscriber ID data extraction**: Specify the regular expression for subscriber ID if you want to extract specific data from the RADIUS accounting messages. The regular expression can contain as many sub-expression groups that you may specify in the next field.
  - **Subexpression**: Choose a number of the sub-expression groups from the drop-down list. If the sub-expression is set to zero, then the text matching the entire regular expression is used for data extraction. If it is non-zero, then the REGEX must contain at least that many sub-expression groups.
  - **Alternate Subscriber ID**: Select the AVP that you want to use as the alternate subscriber ID. For example, the alternate subscriber ID can be used to identify the fixed line or home gateway router. Depending on the AVP you select in this field, the appliance extracts those AVPs from the RADIUS accounting messages. All the sites in the Grid will use the AVP selected in this field as the alternate subscriber ID. Note that the alternate subscriber ID configuration will override the subscriber ID configuration when both the AVPs are available in the RADIUS account messages.  
You can select one of the following AVPs from the drop-down list: APN, Calling-Station-Id, Class, IMEI, IMSI,

MSISDN, and User-Name. Note that for subscribers who have opted for Parental Control service, you can select User-Name as the alternate subscriber ID.

- **REGEX for Alternate Subscriber ID data extraction:** Specify the regular expression for alternate subscriber ID if you want to extract specific data from the RADIUS accounting messages. The regular expression can contain as many sub-expression groups that you may specify in the next field.
- **Subexpression:** Choose a number of the sub-expression groups from the drop-down list. If the sub-expression is set to zero, then the text matching the entire regular expression is used for data extraction. If it is non-zero, then the REGEX must contain at least that many sub-expression groups.
- **Local ID:** Select **LocalID** AVP from the drop-down list to identify subscribers behind a gateway. For example, a home gateway. Each device behind the home gateway network is identified by their local ID/ client ID which is the MAC address of the subscriber device. The appliance extracts the local ID/client ID from the RADIUS accounting messages. Note that for subscribers who have opted for Parental Control service, must select **LocalID** AVP. The home gateway router attaches the MAC address of the host in an EDNS0 option to DNS query request. The DNS Cache Acceleration extracts the EDNS0 option and uses the value to look up for the relevant policy. The client identification information defined in dnsmasq version 2.78 is used to determine the client ID. The MAC address can only be added if the requestor is on the same subnet as the dnsmasq server. Note that the mechanism used to achieve this (an EDNS0 option) is not yet standardized, so this should be considered experimental. Also, note that exposing MAC addresses in this way may have security and privacy implications.
- **REGEX for Local ID data extraction:** Specify the regular expression for local ID if you want to extract specific data from the RADIUS accounting messages. The regular expression can contain as many sub-expression groups that you may specify in the next field.
- **Subexpression:** Choose a number of the sub-expression groups from the drop-down list. If the sub-expression is set to zero, then the text matching the entire regular expression is used for data extraction. If it is non-zero, then the REGEX must contain at least that many sub-expression groups.
- **IPv6 Anchor IP Address:** The ordered list of IPv6 anchor IP address AVP precedence. Note that you cannot modify the IPv6 anchor list, but you can change the order of the list.
- **Subexpression:** Choose a number of the sub-expression groups from the drop-down list. If the sub-expression is set to zero, then the text matching the entire regular expression is used for data extraction. If it is non-zero, then the REGEX must contain at least that many sub-expression groups.
- **NAS Contextual Information:** Select one of the following from the drop-down list: APN, NAS-IP-Address, NAS-IPv6-Address, NAS-Identifier, NAS-Port, or NAS-Port-Identifier.
- **Ancillary Fields:** You can select a list of ordered AVP ancillary fields. Select an ancillary field from the Available column and click the right arrow to move it to the **Selected** column.
- **Interim Accounting Interval:** Specify the time interval in minutes for the RADIUS accounting data to be fully populated in the subscriber collector. The Subscriber Service must be active (green state) before Anycast can serve. The green state is dependent on the value you specify in the Interim Accounting Interval field. That is, the time interval that the subscriber cache has been processing accounting RADIUS messages. RADIUS must send a START or an UPDATE accounting message for every active subscriber during the interim accounting interval. Once this interval passes, the cache is considered fully populated and then Anycast is allowed to serve. If Parental Control is enabled, Anycast needs to wait for the category database to be available before it is allowed to serve. Anycast is allowed to serve only if the Subscriber Service is green. If the Subscriber Service is yellow, it means that either the interim accounting interval has not passed, or that the category database is not available. Note that if you have configured DNS Anycast for the Grid member, then DNS Anycast is disabled for the member during the interim accounting interval.
- **Collect on the MGMT interface only:** Select this checkbox if you want the NAS RADIUS traffic to be accepted over the MGMT interface only.

## Device Handling Performance Optimization for NIOS Subscriber Cache

The NIOS subscriber cache supports various subscribers, including mobile devices, households (CPEs), and devices behind each CPE. The devices behind a CPE include registered devices with policies provisioned or unregistered devices without policies in the Harmony database. Mobile devices and CPEs are classified as provisioned devices.

In NIOS, provisioned devices are not allowed in the Least Recently Used (LRU) classification, but all other devices including unregistered active devices can be allowed in the LRU. The devices in the LRU classification are not subject to garbage collection, and will be deleted as they age due inactivity. Unregistered devices are deleted first based on the LRU classification, to prevent the subscriber's cache from exceeding its limit.

The Harmony database accepts the NIOS modifications for more efficient use of the NIOS cache. Device handling performance optimization optimizes the use of NIOS subscriber cache by updating only the provisioned devices to NIOS thereby reducing the number of devices delivered to the NIOS cache.

## Infoblox Subscriber Parental Control

The Infoblox Subscriber Parental Control provides a mechanism to enable subscribers to manage Internet access and content for their mobility devices, houses, families, or corporations based on the content categories and domains. This helps in restricting users from accessing certain specific content, especially restrict children from accessing specific websites, that are deemed inappropriate. Each subscriber who has opted for the service must be associated with a filtering profile that includes the categories to be blocked for the subscriber session. You can use the pre-defined profiles that address different population segments (such as child, youth, young adults, etc.) or you can create custom profiles for each subscriber in the Infoblox Subscriber Interface. For example, you can define profiles, such as to block children from accessing gambling websites, allow access to educational websites, and monitor access to entertainment websites. You can define profiles for a specific time of the day and for a specific duration. For example, parents can block children from accessing gaming websites from 7.00 AM to 9.00 PM every day. RPZs are used to perform content filtering for the subscribers who have opted for the service. Whenever a subscriber query matches the content of any RPZ, flagged by the blocked category, the traffic is blocked and redirected to the blocking VIP addresses.

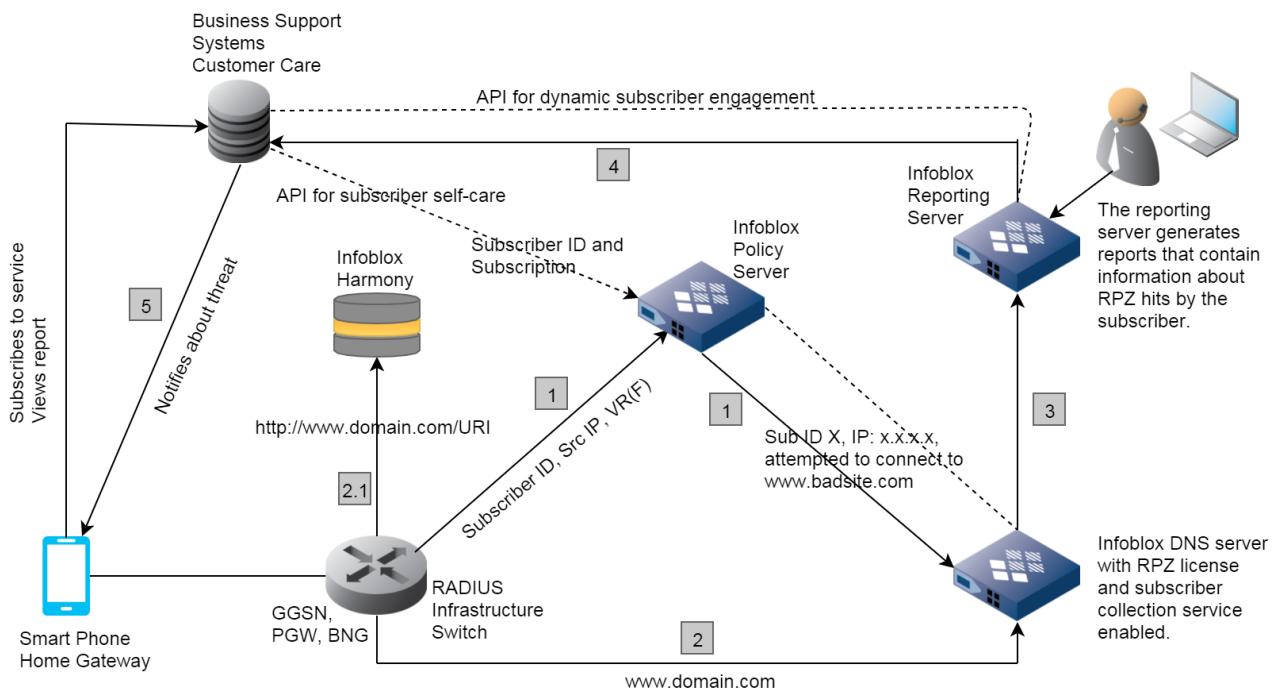
You can also add exceptions for the blocked or allowed categories for each profile or policy in the Infoblox Subscriber Interface. For example, you can block gambling websites but allow casino.com or allow alcohol websites but block vodka.com. You can add a maximum of 10 domains each to the blacklist and whitelist domains for each subscriber. The subscriber query is matched with the blacklist and whitelist domains and appropriate action is taken. If a query is matched with a blacklisted domain, the query is redirected to the blocking server and if a query is matched with a whitelisted domain, the query is resolved normally.

The Infoblox Subscriber Parental Control is currently supported on the following Infoblox appliances: IB-1415, IB-1425, IB-2215, IB-2225, PT-1405, PT-2205, IB-4030-10GE, IB-VM-1405, IB-VM-1415, IB-VM-1425, IB-VM-2205, IB-VM-2225, IB-VM-1425, and IB-FLEX.

As illustrated in the following Infoblox Subscriber Parental Control figure, the DNS server receives RADIUS accounting messages and AVPs (Attribute Value Pairs) from the Infoblox Harmony product. For information about Infoblox Harmony product, refer to Infoblox Harmony documentation. The AVP includes the policy vector that defines the blocked categories and domains for the subscribers who has opted for the service. RPZs perform the filtering of content for these subscribers by applying the parental control policies on incoming subscriber queries. The appliance either allows or blocks the traffic based on the parental control policies. The blocked traffic might also be redirected to the MSP server for evaluation of the traffic. When the traffic is blocked, a blocking page is displayed to the subscriber describing the reason for blocking the traffic. The parental control policies are configured on the Infoblox Harmony product using the Infoblox Subscriber Interface and the policies can have an expiration date.

The NIOS appliance logs all parental control related events, conformed to CEF (Common Event Format), in the syslog. You can get information about the hit when users try to access one of those websites on the blocking list. The reporting server in the Grid generates corresponding reports that contain statistics about parental control related events. For information about monitoring parental control hits by users, see [Monitoring Subscriber Policy Violations](#).

*Infoblox Subscriber Parental Control*



**Note:**  
 2.1: DNS resolves to proxy if blocking is needed or if URI inspection is needed.  
 3: RPZ/PC logs sent to the reporting server for analytics.

## Guidelines for Using Infoblox Subscriber Parental Control

Following are some guidelines to take into consideration when using Infoblox Subscriber Parental Control:

- Infoblox recommends that you configure Infoblox Subscriber Parental Control only on appliances that have a disk size greater than 250 GB.
- You must configure the subscriber collection service and ensure that it is running properly in order to enable Infoblox Subscriber Parental Control.
- The appliance processes DNS queries using the standard DNS query processing, for unknown subscribers.
- You must restart the DNS service after you enable or disable the Subscriber Parental Control.
- To enable Subscriber Parental Control, the Infoblox Harmony product is required. You must configure a DTC (DNS Traffic Control) health monitor to monitor the health of the Infoblox Harmony product.
- Enabling Subscriber Parental Control automatically creates an authoritative zone in the default DNS view with A, AAAA, and CNAME records for Infoblox Harmony and the blocking VIP addresses.
- Deleting a record from the parental control authoritative zone can lead to a service interruption.
- If you make any changes to the parental control policy configuration in the Infoblox Subscriber Interface, it might take up to 15 minutes for the changes to take effect.

## Configuring Infoblox Subscriber Parental Control

To enable and configure Infoblox Subscriber Parental Control on the supported Infoblox appliances, complete the following:

1. Obtain and install a valid Grid-wide or member level **RPZ** license. But for IB-FLEX members, the **RPZ** license is included in the **FLEX Grid Activation** license. For information about licenses license requirements and admin permissions, see [Infoblox Subscriber Insight and Subscriber Policy Enforcement](#). You can also configure a reporting appliance in the Grid to generate reports that contain statistics about RPZ related events.
2. Configure admin permissions so admin users can manage the Infoblox Subscriber Service related tasks. For information about how to configure admin permission, see [About Administrative Permissions](#).

3. Create a subscriber site with at least one Grid member. For information about adding subscriber sites, see [Scaling Using Subscriber Sites](#). It is recommended to add more than one Grid member to the subscriber site for redundancy. You can add a maximum of five Grid members to the subscriber site.
4. Start the subscriber collection service on all the members in the subscriber site, as described in starting and stopping the subscriber collection service, see [Infoblox Subscriber Insight Policy](#).
5. Create RPZs in the default DNS view and specify the order of RPZs. Note that only the default DNS view is supported for configuring RPZs for Subscriber Services. For information about creating RPZs, see [Configuring Local RPZs](#). For information about specifying the order of RPZs and Reordering RPZs, see [Managing RPZs](#). You can create a total of 32 RPZs, out of which the first five RPZs are employed for DNS Firewall and are applied to all DNS lookups. The rest of the RPZs are applied to selective subscribers depending on the policies. The parental control policies received from Infoblox Harmony enables the selection of the RPZs applicable for the subscribers.
6. Enable all the members in the Grid to respond to recursive queries, as described in [Enabling Recursive Queries](#). Add a forwarder and enable **Use Forwarders Only** option in the **Forwarders** tab of the *Member DNS Properties* editor.
7. Enable RPZ logging in the *Member DNS Properties* editor for each member of the site, to ensure that all events related to RPZ are logged to the Syslog. Note that you can also enable logging of queries and responses, but it might significantly affect system performance. For information about how to set DNS logging categories, see [Using a Syslog Server](#).
8. After completing the DNS configuration on the Grid members, start the DNS service on the Grid members. For information about how to start and stop the DNS service, see [Starting and Stopping the DNS Service](#). Note that you ensure to enable IPv6 on the Grid members to support IPv6 subscribers. For information about configuring IPv6 on a Grid Member, see [Understanding DNS for IPv6](#).
9. You can add AVPs that are not available in the list of predefined AVPs. For information about adding AVPS, see [Managing AVPs \(Attribute Value Pairs\)](#).
10. Configure the subscriber ID settings to associate an AVP with the subscriber in the *Subscriber Services Properties* editor, as described in the Configuring Subscriber Services Properties, see [Infoblox Subscriber Insight Policy Enforcement](#).
11. Enable Subscriber Parental Control in the *Subscriber Services Properties* editor, as described in Enabling Subscriber Parental Control below.
12. Configure the blocking VIP addresses and add IP addresses of the policy management server in the **Parental Control** tab of the *Subscriber Site Properties* editor for each subscriber site. For information about Modifying Subscriber Sites, see [Scaling Using Subscriber Sites](#).
13. Define subscriber profiles and subscriber policies on the Infoblox Harmony product using Infoblox Subscriber Interface. For information, refer to Infoblox Harmony documentation.
14. Assign subscriber profile and subscriber policy for each subscriber who has opted for the service in the Infoblox Subscriber Interface. For information, refer to Infoblox Harmony documentation.
15. Publish parental control policies in the Infoblox Subscriber Interface. For information, refer to Infoblox Harmony documentation.

After you set up the Infoblox Subscriber Parental Control, you can monitor parental control related events using predefined reports and the syslog, as described in [Monitoring Subscriber Policy Violations](#).

#### **Note**

If the subscriber site has HA and the HA passive node is the first to upgrade, the data repository connectivity uses the IPv4 protocol for the site members. If you want the data repository to be connected over the IPv6 protocol, you must stop and restart the subscriber service in the upgraded Grid. The subscriber data is lost when the service is stopped and restarted. It is recommended to stop/start the service of each member at a time to synchronize the subscriber cache with the next active member on the same site.

## Enabling Subscriber Parental Control

To enable Subscriber Parental Control, you must ensure that the subscriber collection service is configured and is running properly. After enabling Subscriber Parental Control, you must add Parental Control blocking VIP addresses and add at least one MSP server to the subscriber site. For information about adding Parental Control blocking IP addresses



and MSP addresses to the subscriber site and Modifying Subscriber Sites, see [Scaling Using Subscriber Sites](#). Enabling Subscriber Parental Control, automatically creates an authoritative zone in the default DNS view with A, AAAA, and CNAME records for the MSP and the blocking VIP addresses.

 **Note**

You can disable Subscriber Parental Control only when all the Parental Control blocking VIP addresses and the MSP addresses in all the subscriber sites are removed.

To enable Infoblox Subscriber Parental Control:

1. From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab, expand the Toolbar and click **Subscriber Services Properties**.
2. In the **Parental Control** tab, complete the following:
  - **Enable Parental Control**: Select this checkbox to enable Subscriber Parental Control globally.
  - **Category Information**: Complete the following to add information about the server for the category feed: Contact Infoblox Technical Support for the category feed account information.
    - **User Name**: Enter the logon name of the user.
    - **User Password**: Enter the password of the user.
    - **Server URL**: Enter the URL of the categorization feed server for the category feed.
    - **Update Interval in hours**: You can specify the time interval in hours that determines how often the category information is updated.
  - **Category Proxy**: Complete the following to add information about the proxy server for the category feed.
    - **Proxy URL**: Enter the URL of the proxy server for the category feed.
    - **Proxy User Name**: Enter the logon name of the user.
    - **Proxy User Password**: Enter the password of the user.
  - **Local Zone Name**: Enter the name of a local authoritative zone in the **Name** field.

## Scaling Using Subscriber Sites

You can create subscriber sites and add a Grid member as a collector member and RPZ members to the site in order to scale the number of subscribers that the system can support. The subscriber collector caches the subscriber data received from the NAS gateways and parental control policies from the Infoblox Harmony product. The RPZ members use the cached subscriber data and the policies to resolve DNS queries. You can add a maximum of five Grid members to the subscriber site. Note that one Grid member can serve only one subscriber site. The subscriber identity information cached in the subscriber cache is replicated between the Grid members in the subscriber site.

You can configure the NAT port as an IPSD, where the subscribers first deterministic NAT port is used as IPSD, to distinguish from other subscribers using the same IP address. The NAT algorithms use the port range allocation where the first usable port for the subscriber is provided in a RADIUS accounting AVP. It also supports the first port in the range as a discriminator between subscribers using the same IP address. In a strict NAT configuration, where only the NATed subscribers are allowed, the value of the AVP **Deterministic NAT port** must be a non zero value in the RADIUS accounting message. The ports from 1-1023 (inclusive) are reserved in a deterministic port configuration.

You can manage the subscriber sites as discussed in the following sections:

- [Adding Subscriber Sites](#)
- [Bypassing Subscriber Secure Policy for Allowed Lists](#)
- [Modifying Subscriber Sites](#)
- [Deleting Subscriber Sites](#)
- [Viewing Subscriber Sites](#)
- [Viewing NAS Gateway Message Rates](#)

Additionally, you can import subscriber site data by using the **CSV Import** option on the Toolbar.

## Limitations using NAT port as IPSD

- No acceleration support using SNIC appliances, however, it is supported for all appliances including vDCA acceleration.
- You need to restart the DNS service.
- IPSD is a global configuration that applies to all Sites. IPSD may be set from CSV, WAPI, and CLI to a different AVP.
- A performance penalty for Dynamic subscribers (without Deterministic-Nat-Port AVP) in Deterministic Sites (Sites with block\_size configuration) requires two lookups.
- NAT port can be configured as IPSD only if the subscriber services properties is set to **Deterministic-Nat-Port** and block size must be greater than zero.
- The site block size must be the same as the deployment CGNAT block size configuration.
- Changing the site block size will initialize the state of the subscriber collection.
- Static default network policies in a strict NAT configuration (Allow NATed Subscribers only) will not resolve at the DCA.

## Adding Subscriber Sites

To add a subscriber site, complete the following:

1. From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab -> **Subscriber Sites** tab, click the Add icon.
2. In the *Add Subscriber Site* wizard, complete the following:
  - a. **Name:** Enter the name of the subscriber site.
  - b. **Maximum Subscribers:** Specify the maximum number of subscribers for the subscriber site. This represents the overall size of the subscriber cache. You can enter a value between 10000 to 10000000.
  - c. **Comment:** You can enter additional information about the subscriber site.
  - d. **Members:** In the Members table, click the Add icon to add Grid members to the site. If there are multiple members, the *Member Selector* dialog box is displayed, from which you can select a member. Click the required member name in the dialog box. You can also delete a member from the list.  
Note that a Grid member can support only one subscriber site.
  - e. **Deterministic NAT Block Size:** The block size specifies the number of ports made available for each incoming subscriber address. In a deterministic NAT, zero means not using NAT. The value can be any number from 0 to 64512. The block size configuration is not allowed to change unless the global (subscriber service properties) IPSD is set to Deterministic-NAT-Port.
  - f. **First port:** The value of the first usable port for the subscriber. The first usable port will have a default value of 1024, and the value can be any number from 1024 to 65535, both inclusive.
  - g. **Allow NATed Subscribers only:** Select this option to restrict only NATed subscribers (Subscribers with IPSD). Here the IP address and port block allocations are made dynamically for the subscriber instance and the IPSD of the first port is assigned to the subscriber port block. For example, if the block size is 8 for the site, then the IPSD must be set to 1024, 2032, 3040, etc.  
**Stop the anycast service when the subscriber service is in the interim state:** Select this option to stop the anycast service from running when the subscriber service is in the interim state. By default, this option is selected.
3. Click **Next** to configure NAS gateways for the subscriber site. Complete the following
  - a. **Listen on RADIUS port number:** Enter the UDP port number that the collector member uses to collect accounting information from the NAS gateway. You can enter an integer from 1 to 65535. The default is 1813.
  - b. **NAS Gateways:** You must add at least one NAS gateway to the subscriber site in order to start the subscriber collection service. You can add up to 20 NAS gateways. Click the Add icon and complete the following to add a NAS gateway:
    - i. **Name:** Enter the name of the NAS gateway.
    - ii. **IP Address:** Enter the IP address of the NAS gateway.
    - iii. **Shared Secret:** Enter a shared secret that can be used to authenticate the communication between the RADIUS accounting server and the collector member. This shared secret must match the one you entered on the RADIUS server.
    - iv. **Confirm Shared Secret:** Enter the shared secret again.



- v. **Send Protocol Acknowledgment:** Select this checkbox to send an acknowledgment to the client when the collector member receives accounting information from the NAS gateway.
  - vi. **Comment:** Enter additional information about the NAS gateway.
  - vii. Click **Add** to add the NAS gateway.  
You can select a NAS gateway configuration and click the Edit icon to modify it or click the Delete icon to delete it.
4. *This step is required only if Infoblox Subscriber Parental Control is enabled.* For information about enabling Parental Control, see [Infoblox Subscriber Parental Control](#). Click **Next** to configure the parental control blocking IP addresses. Complete the following:
- a. **Content Proxy Addresses:** You can add IP addresses of the Infoblox Harmony product. The appliance will forward the subscriber session to Infoblox Harmony for in-line processing of the subscriber session, depending on the policies. Click the Add icon. Grid Manager adds a row to the **Content Proxy Addresses** table. It is recommended that you enter two addresses in this field. The first address is considered the primary address and the second address is considered the secondary address. If you enter only one address, the same address is considered the primary and secondary address. Click the row and enter the IP address in the **Address** field. To delete an IP address, select the checkbox and then click the Delete icon.
  - b. **Enforce the global proxy list:** Select this checkbox if you want to proxy the traffic to the MSP (Multi-Services Proxy) server. If you select this checkbox, and have categorized the queried domains in the incoming traffic to the global proxy list using the `set pc_domain add` command (category 104), then the query resolves to an MSP virtual IP address and NIOS generates a "synthetic resolution". This checkbox is disabled by default, and you must configure **Content Proxy Addresses** to enable it. If you do not select this checkbox, then the query resolves normally.  
Notes:
    - If you have configured queries to specific domains (categorized to 104) to be proxied to the MSP server and enabled the **Enforce the global proxy list** option, queries to these domains are proxied if subscriber secure policies with the NXDOMAIN rule are not set.
    - If you want to enable and run DNS over TLS, DNS over HTTPS, and Parental Control features simultaneously on a member, ensure that the appliance meets the base memory configuration requirements defined in [Configuration Requirements](#). If you try to run these features when the required memory configuration is not available, all of these features will be disabled.
  - c. **Additional Blocking Servers:** Besides the IP addresses you specify in the Parental Control Blocking IP Addresses fields, you can specify additional IP addresses that will act as blocking servers for the blocking policies you defined when [configuring blocking server policies](#). Click the Add icon. Grid Manager adds a row to the **Additional Blocking Servers** table. Click the row and select a blocking policy. In the **Address** field, enter the IP address of the blocking server that will contain the selected blocking policy. To delete an IP address, select the checkbox and then click the Delete icon.
  - d. **Parental Control Blocking IP Addresses:** You can configure two sets of IPv4 and IPv6 addresses that are used as blocking VIP addresses. The parental control subscribers are redirected to the following blocking IP addresses whenever the domain queried by the subscriber is blocked based in the subscriber parental control policy.  
Complete the following:
    - i. **IPv4 Address (primary):** Enter the primary blocking IPv4 address.
    - ii. **IPv4 Address (secondary):** Enter the secondary blocking IPv4 address.
    - iii. **IPv6 Address (primary):** Enter the primary blocking IPv6 address.
    - iv. **IPv6 Address (secondary):** Enter the secondary blocking IPv6 address.
  - e. **Policy Management Addresses:** You can add IP addresses of the policy management servers to which the appliance sends APIs about the expired parental control policies. Click the Add icon. Grid Manager adds a row to the **Policy Management Addresses** table. Click the row and enter the IP address in the **Address** field. To delete an IP address, select the checkbox and then click the Delete icon
5. Save the configuration, or click **Next** to continue to the next step where you define extensible attributes as described in [Managing Extensible Attributes](#).

## Bypassing Subscriber Secure Policy for Allowed Lists

You can choose to let subscriber specific allowed domains to take priority over category-based policies, security policies and blocklist entries. Subscriber specific blocked domains take priority over category-based policies.

To bypass subscriber secure policy for allowed lists:

1. Navigate to the **Data Management > DNS > Subscriber Services Deployment > Subscriber Sites** tab.
2. Click the Edit icon to edit an existing subscriber site.
3. In the *Subscriber Site Properties* wizard, go to the **General > Advanced** tab.
  - **Enable Subscriber Secure Policy Bypass for Allowed list:** Select this checkbox to enable NIOS to generate a normal response for all domains in a subscriber's allow list. The allowed domains override RPZ rules if any (for example, NXDOMAIN), and categorize policy rules for the subscriber. This enables subscribers to override all policies for a specific domain. The subscribers allow list domain is cached in DNS Cache Acceleration and subsequent queries are answered by DNS Cache Acceleration. Note that the vDCA ( virtual DNS Cache Acceleration) allows only 5% of its total number of subscribers to have allow block lists. The maximum number of allowed and blocked domains is 15. Domains in a subscriber's blocked list take priority over category-based policies subject to RPZ rules. This ensures that the RPZ global allow list is always enforced. That is, a subscriber cannot override a domain in the global allow list.
  - **Set Global Allow List RPZ index range ( 0 to 30):** Select this checkbox to specify an RPZ index value between 0 to 30. A domain is added to the RPZ specified as a passthru RPZ rule, and that domain is added as a global allowed list. This global allowed list is applicable to all subscribers. You can use both the **Enable Subscriber Secure Policy Bypass for Allowed list** and the **Set Global Allow List RPZ index range** options at the same time. Or you can choose to use the options independently irrespective of whether **Enable Subscriber Secure Policy Bypass for Allowed list** is enabled or disabled.

## Modifying Subscriber Sites

To modify a subscriber site, complete the following:

1. From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab -> **Subscriber Sites** tab, click the **Action** icon next to the subscriber site name and select **Edit** from the menu.
  2. The *Subscriber Site Properties* editor provides the following tabs from which you can modify data:
    - a. In the **General** tab, you can modify the information you previously entered through the *Add Subscriber Site* wizard.

Note that If the Grid Manager is IPv4-only and Grid members are IPv6 or dual-stack, modifying subscriber sites may not clear the subscriber cache. Import and export of subscriber data may also not work.
    - b. In the **NAS Gateways** tab, you can edit the NAS gateways configured for the subscriber site, as described in the *Adding Subscriber Sites* section.

Note that If you make any changes to the NAS gateway configuration, the subscriber collector will automatically restart within 30 seconds. However, the subscriber data collected in the subscriber cache is not affected by the NAS gateway configuration changes.
    - c. If Subscriber Parental Control is enabled, the **Parental Control** tab is displayed. You can modify the information in the **Parental Control** tab, as described in the section Adding Subscriber Sites above.
      - i. **Enable DCA subscriber Query count logging:** Select this checkbox to allow the DCA to generate subscriber logs and to record query counts greater than or equal to zero for subscriber query count updates and deletions. These logs are generated for deletions even when the query count is equal to zero. By default, this option is disabled.
      - ii. **Enable DCA subscriber Allowed & Blocked list support:** Select this checkbox to support the blocked and allowed list of subscribers. This option is disabled by default. Once the domain is cached, the blocked lists are provided by DCA. Domains in the allowed list are transferred to BIND. There are several members on the site, but the memory requirement is 32GB or higher for all the vDCA capable members. You must manually restart NIOS after selecting this checkbox for the support to be successful.
- Notes:

- The allowed and blocked listing feature allows you to specify all possible top-level domains, (for example, <http://linkedin.com> , [linkedin.co.uk](http://linkedin.co.uk)) for well-known names. If a dotless name such as "facebook" is in the allowed list or blocked list and the qname is facebook.<suffix>, then:
    - If the suffix is a top-level domain (example "xxxxyy"), the two are matched regardless of whether "xxxxyy" is registered or not in the worldwide DNS. Example:  
facebook == <http://facebook.com>  
facebook == facebook.xxxxxy
    - If the suffix is not a top-level domain (example "xxx.yyy"), whether the two are matched or not depends on whether "xxx.yyy" is registered and present in the public\_suffix\_list.dat on the appliance or not. Example:  
facebook == [facebook.co.uk](http://facebook.co.uk)  
facebook != facebook.xxx.yyy
  - If you remove a policy from the Proxy-All allow list, wait for the time to leave (TTL) setting that is configured in DNS Cache Acceleration to expire, for the subscriber policy to work correctly.
  - If a zvelo category database update failure occurs for three consecutive days:
    - Grid Manager displays a yellow background with the "Please correct the download credentials or the proxy configuration to get the latest database updates" message and the member status is displayed as "Domain category db is not latest" in the **Grid Manager >Subscriber Collection > Services> Service Status** column.
    - A new SNMP trap is sent with the message "Domain category db is not latest". Additionally, if email notifications are configured, an email is sent to the configured email address with the "Domain category db is not latest" message.
    - Post this event, if the zvelo download is successful, a new SNMP clear trap is sent, and an email with the "zvelo SNMP Clear Trap" message is also sent. The **Service Status** column is on green background will be displayed in Green with the "Subscriber Collection Service is working" message.
  - If a zvelo category database update failure occurs for more than 60 days:
    - Grid Manager displays a red background with the "zvelo database expired. Subscriber secure queries will be fail-open" message and the member status is displayed as "zvelo db has expired" in the **Grid Manager > Subscriber Collection > Services > Service Status** column.
    - A new SNMP trap is sent with the "zvelo db has expired" message. Additionally, if email notifications are configured, an email is sent to the configured email address, with the "zvelo db has expired" message.
    - Post this event, if the zvelo download is successful, a new SNMP clear trap is sent, and an email with the "zvelo SNMP Clear Trap" message is also sent. The **Service Status** column is on green background will be displayed in Green with the "Subscriber Collection Service is working" message.
  - If you have not downloaded the zvelo database earlier:
    - If the zvelo database download fails for 3 consecutive days, a critical SNMP trap is sent, and "Domain category db is not latest" is displayed as the member status instead of "Category information data is unavailable".
    - If the zvelo database download fails for 60 consecutive days, a critical SNMP trap is sent, and "Domain category db is expired" is displayed as the member status instead of "Category information data is unavailable".
  - d. You can enter or edit information in the **Extensible Attributes** tab, as described in [Managing Extensible Attributes](#).
  - e. You can export subscriber site data into a CSV file by selecting the **Export** option. For more information, see [Importing and Exporting Data using CSV Import](#).
3. Save the configuration.

## Deleting Subscriber Sites

To delete a subscriber site, complete the following:

From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab -> **Subscriber Sites** tab, click the Action icon next to the subscriber site name and select **Delete** from the menu.

In the *Delete Confirmation (Subscriber Site)* dialog box, click **Yes**.

## Viewing Subscriber Sites

To view subscriber sites, complete the following:

From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab -> **Subscriber Sites** tab.

Grid Manager displays the following information for each subscriber site:

1. **Actions:** Click the Action icon next to a selected subscriber site and choose from the following:
  - a. **Edit:** Modify certain general properties.
  - b. **Delete:** You can delete the subscriber site.
  - c. **Extensible Attributes:** Add or modify extensible attributes.
  - d. **View NAS Gateway Message Rates:** Displays the message rates of the NAS gateways configured for the subscriber site.
2. **Name:** The name of the subscriber site.
3. **Comment:** Information about the subscriber site.
4. **Site:** Displays values that were entered for this predefined attribute.

You can also perform the following:

1. Edit the subscriber site information. Select the subscriber site, and then click the Edit icon.
2. Delete a subscriber site. Select the subscriber site, and then click the Delete icon.
3. Export the list of subscriber sites. Click the Export icon.
4. Print the list of subscriber sites. Click the Print icon.

Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches.

Create a quick filter to save frequently used filter criteria:

1. In the filter section, click **Show Filter** and define filter criteria for the quick filter.
2. Click **Save** and complete the configuration in the *Save Quick Filter* dialog box.

The appliance adds the quick filter to the quick filter drop-down list in the panel. Note that global filters are prefixed with [G], local filters with [L], and system filters with [S].

Sort the subscriber sites in ascending or descending order by column.

## Viewing NAS Gateway Message Rates

You can view the NAS gateways (accounting log servers) configured for the subscriber site and the message rate for each NAS gateway.

To view the NAS gateway message rates, complete the following:

From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab -> **Subscriber Sites** tab.

In the **Subscriber Sites** tab, click the **Action** icon next to the respective subscriber site and select **View NAS Gateway Message Rates** from the list.

The *NAS Gateway Message Rates* dialog box displays the following information for the selected subscriber site:

**Name:** The name of the NAS gateway.

**IP Address:** The IP address of the NAS gateway.

**Message Rate:** The message rate of the NAS gateway.

## Configuring Blocking Server Policies

In addition to the default blocking servers that you specify when adding subscriber sites, you can also create service policies that can be associated with specific servers. These servers act as blocking servers and any traffic or web pages that conform to the service policies that you create are blocked and redirected to the blocking VIP addresses.

### Adding a Blocking Server Policy

To create or add a blocking server policy:

1. From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab, expand the vertical Toolbar and click **Blocking Server Associations**.
2. Click the Add icon.
3. In the **Name** field, specify a name for the policy.
4. In the **Parental Control Policy** field, specify the hexadecimal value of the parental control policy that you want to add. For a list of the different types of parental control policies and their hexadecimal values, refer to the supplemental documentation provided by Infoblox or contact your Infoblox representative.
5. Click the **Add** button.

The policy that you created is displayed in the table. You can then associate IP addresses with the policy that you created. For more information, see [Scaling Using Subscriber Sites](#).

### Editing a Blocking Server Policy

To edit an existing blocking server policy:

1. From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab, expand the vertical Toolbar and click **Blocking Server Associations**.
2. Select the policy that you want to edit and click the Edit icon.
3. Edit either the name of the policy or its hexadecimal value.
4. Click **Save**.

### Deleting a Blocking Server Policy

To delete an existing blocking server policy:

1. From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab, expand the vertical Toolbar and click **Blocking Server Associations**.
2. Select the policy that you want to delete and click the Delete icon.
3. Click **Save**.

## Managing AVPs (Attribute Value Pairs)

The subscriber site receives the subscriber security policies and parental control policies as AVPs from the RADIUS server. You can create AVPs that are not available in the list of predefined AVPs and manage the AVPs in the following ways:

- [Adding AVPs](#)
- [Modifying AVPs](#)
- [Deleting AVPs](#)
- [Viewing AVPs](#)

- Create an user-defined AVP, as described in Adding AVPs below.
- View the list of AVPs, described in Viewing AVPs below.
- Modify the properties of an user-defined AVP, as described in Modifying AVPs below.
- Delete an user-defined AVP, as described in Deleting AVPs below.

 **Note**

The configured AVPs are applicable for all the subscriber sites.

## Adding AVPs

To add an user-defined AVP:

1. From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab -> **AVPs** tab, and then click the Add icon.
2. In the *Add AVP* wizard, complete the following:
  - **Name:** Enter the name of the AVP (Attribute Value Pair).
  - **User Defined:** Indicates whether the AVP is an user-defined AVP.
  - **Type:** Enter the AVP type as defined in RFC 2865 and RFC 2866. If the AVP type is vendor specific, you can enter the value 26.
  - **Vendor ID:** Specify the vendor ID, if the AVP ID is vendor specific.
  - **Vendor Type:** Specify the vendor type, if the AVP type is vendor specific.
  - **Value Type:** Select the type of value from the drop-down list:
  - **Comment:** You can enter additional information about the AVP.
  - **Make this AVP available as an option in:** Select one of the following:
    - **All fields:** Select this if you want the AVP to be available in all lists in *Subscriber Secure Properties* editor.
    - **Only to these fields:** Select this if you want the AVP to be available only in specific fields in the *Subscriber Secure Properties* editor. Click the Add icon in the table. The appliance adds a row to the table. Select the row and select the field name from the drop-down list.
3. Save the configuration.

## Modifying AVPs

To modify an user-defined AVP:

1. From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab -> **AVPs** tab, click the **Action** icon next to the user-defined AVP name and select **Edit** from the menu.
2. In the *AVP Properties* editor, modify the AVP information, as described in the previous section, Adding AVPs.
3. Save the configuration.

Note that you cannot modify the properties of a predefined AVP. You can only view the information in the *AVP Properties* editor.

## Deleting AVPs

You can delete only user-defined AVPs.

To remove an user-defined AVP:

1. From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab -> **AVPs** tab, click the **Action** icon next to the user-defined AVP name and select **Delete** from the menu.
2. In the *Delete Confirmation (AVP)* dialog box, click **Yes**.

## Viewing AVPs

To view all the predefined and user-defined AVPs:

1. From the **Data Management** tab -> **DNS** tab -> **Subscriber Services Deployment** tab -> click the **AVPs** tab.
2. Grid Manager displays the following information for each AVP:
  - **Actions:** Click the Action icon next to a selected user-defined AVP and choose from the following:
    - **Edit:** Modify certain general properties.
    - **Delete:** You can delete only user-defined AVP.
  - **Name:** The name of the AVP.
  - **User Defined:** Indicates whether the AVP is predefined or user-defined AVP.
  - **Type:** The type of AVP.
  - **Vendor ID:** The vendor ID.
  - **Vendor Type:** The vendor type.
  - **Value Type:** Indicates the value type. Example: string, integer, IPv4 address, IPv6 address, byte, etc.
  - **Comment:** Information about the AVP.
  - **Restricted To:** Indicates whether the AVP is restricted to domains.

You can also do the following:

- Edit the AVP information.
  - Select the AVP, and then click the Edit icon.
- Delete an user-defined AVP.
  - Select the user-defined AVP, and then click the Delete icon.
- Export the list of AVPs.
  - Click the Export icon.
- Print the list of AVPs.
  - Click the Print icon.
- Use filters and the **Go to** function to narrow down the list. With the autocomplete feature, you can just enter the first few characters of an object name in the **Go to** field and select the object from the possible matches. This is not applicable for **Restricted To** column.
- Create a quick filter to save frequently used filter criteria:
  - In the filter section, click **Show Filter** and define filter criteria for the quick filter.
  - Click **Save** and complete the configuration in the *Save Quick Filter* dialog box.
- The appliance adds the quick filter to the quick filter drop-down list in the panel. Note that global filters are prefixed with [G], local filters with [L], and system filters with [S].
- Sort the AVPs in ascending or descending order by column. This is not applicable for **Restricted To** column.

## Monitoring Subscriber Policy Violations

You can monitor RPZ hits based on subscriber security policies and parental control policies through the following:

- [Monitoring through Syslog](#)
- [Handling Splunk REST API Request](#)
- Syslog, as described in [Monitoring through Syslog](#) below.
- [Detailed RPZ Violations by Subscriber ID](#) report, as described in [Detailed RPZ Violations by Subscriber ID](#) below.

## Monitoring through Syslog

To receive information about RPZ hits based on subscriber security policies and parental control policies in the syslog, make sure that you enable the **RPZ** option in the **Logging** tab of the Grid DNS Properties editor or Member Properties editor. For information about configuring logging properties, setting DNS Logging Categories, see [Using a Syslog Server](#). Once the RPZ option is enabled, the appliance logs RPZ hits based on subscriber security policies and parental control policies in CEF (Common Event Format) in the syslog. For information about how to configure the syslog server, see [Using a Syslog Server](#).

Following is a sample RPZ hit log message:

```
CEF:0|Infoblox|NIOs|8.2.0-359884|RPZ-QNAME|PASSTHRU|7|app=DNS
dst=10.35.41.18 src=10.32.1.145 spt=52100 view=_default qtype=A msg="rpz
QNAME PASSTHRU rewrite child.com [A] via child.com.bit6subscribers" IPSD=N/A
Acct-Session-Id=29de847acde415ab User-Name=john NAS-IP-Address=10.36.120.10
MSISDN=9956182386 Subscriber-Secure-
Policy=0000507f CAT=0x000000000000000000000000000008000
```

Each log message contains the following information:

- **Infoblox|NIOs |x.x.x:** Indicates the Infoblox product, and `x.x.x` represents the NIOs version.
- The string following the NIOs version is a hard-coded constant. In this example, it is RPZ QNAME.
- The hard-coded constant is followed by mitigation action. In this example, it is PASSTHRU.
- The number following the mitigation action is the threat severity level. The following numbers indicate the severity levels:
  - **8 = Critical**
  - **7 = Major**
  - **6 = Warning**
  - **4 = Informational**
- **app:** DNS
- **dst:** Destination IP address.
- **src:** Source IP address.
- **spt:** Source port.
- **view:** DNS view.
- **qtype:** Query type.
- **msg:** RPZ rule.
- **IPSD:** IP space discriminator.
- **Acct-Session-Id:** Session ID.
- **User-Name:** Username of the subscriber.
- **NAS-IP-Address:** NAS IP address.
- **MSISDN:** The mobile phone number of the subscriber.
- **Subscriber-Secure-Policy:** Subscriber Secure Policy.
- **CAT:** The category bits that match the query name. This indicates the reason for blocking the domain listed in the event.

To view RPZ violation by subscriber related log messages:

1. From the **Administration** tab, select the **Logs** tab -> **Syslog** tab.
2. From the drop-down list at the upper right corner, select the Grid member on which you want to view the syslog.
3. From the **Quick Filter** drop-down list, select **RPZ Incident Logs** to view RPZ violation by subscriber related events. To narrow down the system messages you want to view, click **Show Filter** and then select the filters you want to use. For information about how to use filters, see [Finding and Restoring Data](#).

## Handling Splunk REST API Request

You can retrieve pre-defined reports by sending Splunk REST API requests. Before you send any Splunk REST API request, ensure that reporting service is enabled on the Grid. To accept the Splunk REST API requests, the apache httpd process is started on the reporting member. The Splunk REST API request process uses TCP port 9185 and you can use NIOs user credentials to send Splunk REST API requests. For information about Splunk REST API, see <http://docs.splunk.com/Documentation/Splunk/latest/RESTUM/RESTusing>.

Note:



- Only searches(/search/jobs/), saved searches(/saved/searches/) and auth login(\* /auth/login) are supported and any other Splunk REST API requests are not supported.
- Splunk REST API requests cannot be sent by remote users.

Following are the supported Splunk REST API requests:

authentication /URI	Summary	GET	PUT	POST	DELETE
auth/login	Access control Provide user authentication			Yes	
saved /URI	Summary	GET	PUT	POST	DELETE
saved/searches	Search Manage saved search configuration	Yes		Yes	
saved/searches/{name}	Search Manage specific saved search	Yes		Yes	Yes
saved/searches/{name}/acknowledge	Search Manage saved search alerts			Yes	
saved/searches/{name}/dispatch	Search Dispatch saved search			Yes	
saved/searches/{name}/history	Search Access saved search job history	Yes			
saved/searches/{name}/reschedule	Search Manage saved search job schedules			Yes	
saved/searches/{name}/scheduled_times	Search Access saved search scheduled time	Yes			
saved/searches/{name}/suppress	Search Access saved search alert state	Yes			
search /URI	Summary	GET	PUT	POST	DELETE
search/jobs	Search Manage search jobs	Yes		Yes	
search/jobs/{search_id}	Search Manage specific search job	Yes		Yes	Yes
search/jobs/{search_id}/control	Search Execute job control command for a specific search			Yes	
search/jobs/{search_id}/events	Search Access events for a specific search	Yes			

search /URI	Summary	GET	PUT	POST	DELETE
search/jobs/{search_id}/results	Search Access results of a specific search	Yes			
search/jobs/{search_id}/results_preview	Search Access preview results for a specific search	Yes			
search/jobs/{search_id}/search.log	Search Access search.log file for a specific search	Yes			
search/jobs/{search_id}/summary	Search Access getFieldsAndStats output of so-far-read events	Yes			
search/jobs/{search_id}/timeline	Search Access event distribution over time	Yes			
search/jobs/export	Search Stream search results	Yes		Yes	

Samples of Splunk REST API requests:

- To search for subscriber IDs:

```
curl -k -u splunk-api-usr:"tru[kl0ad" https://10.61.41.36:9185/services/search/jobs/ -d search="search source=ib:dns:query:top_rpz_hit index=ib_dns| stats count by SUB_VAL"
```

The response is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<response>
  <sid>1516833486.5665</sid>
</response>
```

To get the job status:

```
curl -k -u splunk-api-usr:"tru[kl0ad" https://10.61.41.36:9185/services/search/jobs/1516833486.5665 | grep "isDone"

% Total % Received % Xferd Average Speed Time Time Time Current
 Dload Upload Total Spent Left Speed
100 14702 100 14702 0 0 14702 0 0:00:01 --:--:-- 0:00:01 358k
<s:key name="isDone">1</s:key>
```

To get the job output in CSV format:

```
curl -k -u splunk-api-usr:"tru[kl0ad" https://10.61.41.36:9185/services/search/jobs/1516833486.5665/results --get -d output_mode=csv
```

```
"SUB_VAL",count
"+123 (12) 123-0001",2049
"+123 (12) 123-0002",1978
"+123 (12) 123-0003",2105
"+333 (33) 333-3333",121
"+555 (55) 555-5555",3
"+911 (12) 123-0066",1916
"+911 (55) 555-5555",114

16661230001,11447
16661230002,15464
16661230003,15226
911121230066,11245
```

- To search for RPZ event by a subscriber:

```
curl -k -u splunk-api-usr:"tru[kl0ad" https://10.61.41.36:9185/services/
search/jobs/ -d search='search source=ib:dns:query:top_rpz_hit index=ib_dns
SUB_VAL=16661230002 earliest=-4h latest=-1h | transaction
fields="DOMAIN_NAME" maxspan=1h | eval
MITIGATION_ACTION=case(MITIGATION_ACTION
== "PT", "Passthru", MITIGATION_ACTION == "NX", "Block (No Such Domain)",
MITIGATION_ACTION == "ND", "Block (No Data)", MITIGATION_ACTION == "SB",
"Substitute", MITIGATION_ACTION == "A1", "Substitute (A)", MITIGATION_ACTION
== "A4", "Substitute (AAAA)", MITIGATION_ACTION
== "AA", "Substitute (A/AAAA)", MITIGATION_ACTION == "DN", "Substitute
(Domain Name)", MITIGATION_ACTION == "ER", "None") | table TIMESTAMP
DOMAIN_NAME MITIGATION_ACTION'
```

The response is as follows:

```
<?xml version="1.0" encoding="UTF-8"?>
<response>
  <sid>1516834187.5702</sid>
</response>
```

To get the job status:

```
curl -k -u splunk-api-usr:"tru[kl0ad" https://10.61.41.36:9185/services/
search/jobs/1516834187.5702 | grep "isDone"

<s:key name="isDone">1</s:key>
```

To get the job output in csv format:

```
curl -k -u splunk-api-usr:"tru[kl0ad" https://10.61.41.36:9185/services/search/jobs/1516834187.5702/results --get -d output_mode=csv
```

```
TIMESTAMP,"DOMAIN_NAME","MITIGATION_ACTION"
```

```
"2018-01-24 12:46:07 2018-01-24 12:56:07 2018-01-24 13:06:07 2018-01-24 13:16:07 2018-01-24 13:26:07 2018-01-24 13:36:07 2018-01-24 13:46:07","wpad.tme.infoblox.com","Block (No Such Domain)"
```

```
"2018-01-24 11:36:06 2018-01-24 11:46:06 2018-01-24 11:56:06 2018-01-24 12:06:06 2018-01-24 12:16:06 2018-01-24 12:26:06 2018-01-24 12:36:06","wpad.tme.infoblox.com","Block (No Such Domain)"
```

```
"2018-01-24 10:57:41 2018-01-24 11:07:41 2018-01-24 11:17:41 2018-01-24 11:27:41 2018-01-24 11:34:51 2018-01-24 11:35:06","wpad.tme.infoblox.com","Block (No Such Domain)"
```

## VLAN Management

NIOS allows you to track the VLAN usage in your network thereby allowing you to compare an assigned VLAN with VLANs discovered by Network Insight. You can then generate inventory and conflict reports based on this data. You can:

- Create VLAN resource pools to use in the future
- Track the VLANs that are in use and those that are not in use
- Organize VLAN views and ranges by organization or group
- Assign VLANs to defined networks
- View conflicts if any between assigned and discovered VLANs

To configure VLAN management, perform these steps in the following order:

1. [Create a VLAN view.](#)
2. [Create one or multiple VLAN ranges within a VLAN view.](#)
3. [Create VLAN objects that you can group into a range or view.](#)
4. [Assign a VLAN object to a network.](#)
5. [Generate inventory and conflict reports to view conflicts between an assigned VLAN object and a VLAN object discovered by Network Insight.](#)

You can use the [global search](#) function (**Advanced** tab) to search for VLAN views, ranges, and objects. You can also perform a CSV import or export of the VLAN data. For more information, see [Importing and Exporting Data using CSV Import](#).

All VLAN configuration operations can be scheduled and are subject to administrator approval. For more information, see [Configuring Approval Workflows](#).

For information about permissions related to VLAN objects, see [Administrative Permissions for VLAN Management](#).

## Configuring VLAN Views

You can create a VLAN view and group VLAN objects and VLAN ranges into that VLAN view. A VLAN view allows you to view grouped VLANs in an hierarchical format.

A VLAN view named **default** is automatically created for all NIOS installations starting from NIOS 8.4 and upgrades from earlier versions.

## Viewing a VLAN View

To view the list of VLAN views, perform the following steps:

1. From the **Data Management** tab, click the **VLANs** tab. A list of all the VLAN views is displayed. The following details of the VLAN view are displayed:
  - **Name:** Name of the VLAN view
  - **Start VLAN ID:** Integer from which the VLAN ID can start.
  - **End VLAN ID:** Integer till which the VLAN ID can end.
  - **Comment:** Information for the VLAN view if any.
  - **Allow Range Overlapping:** Whether VLAN ranges with overlapping VLAN IDs are allowed in the VLAN view. For information about overlapping VLAN IDs, see [Allow VLAN Range Overlapping](#) below.
  - **Site:** Extensible attribute for the VLAN view.
2. Click a VLAN view name to view all the VLAN objects and VLAN ranges that belong to the VLAN view.

## Adding a VLAN View

To add a VLAN view:

1. From the **Data Management** tab, click the **VLANs** tab.
2. From the Toolbar, click **Add-> VLAN View**. The *Add VLAN Wizard* dialog box is displayed.
3. In the **VLAN View Name** field, enter a name for the VLAN view.
4. In the **Start VLAN ID** field, enter the minimum ID that can be used by a child VLAN range or VLAN object. VLANs can acquire an ID that is in between the values you enter in the **Start VLAN ID** and **End VLAN ID** fields. The VLAN ID must be greater than or equal to 1 and less than or equal to 4094. A VLAN ID is unique. No VLAN object can share the same VLAN ID.
5. In the **End VLAN ID** field, enter the maximum ID that can be used by a child VLAN range or VLAN object. The VLAN ID must be greater than or equal to 1 and less than or equal to 4094. A VLAN ID is unique. No VLAN object can share the same VLAN ID.
6. Select the **Pre-create VLANs** checkbox if you want NIOS to automatically create VLANs. When you select this checkbox, the **VLAN Name Prefix** field becomes enabled and NIOS creates VLANs between the range you specified in the **Start VLAN ID** and **End VLAN ID** fields. You cannot edit the **Pre-create VLANs** checkbox.
7. In the **VLAN Name Prefix** field, enter a prefix that will be added to the VLAN name. The VLAN name is constructed using the following pattern: `<VLAN Prefix Name> + <VLAN ID>` (the VLAN ID is optionally prefixed with zeros). VLANs are automatically created with the prefix name and VLAN ID. You cannot edit the **VLAN Name Prefix** field.
8. Select the **Allow VLAN Range Overlapping** checkbox if you want to allow different VLAN ranges in the same VLAN view to have overlapping start and end VLAN IDs. Overlapping VLAN ranges will have unique VLAN IDs. For example, if you create a VLAN view with two ranges from 1 to 15 and from 2 to 18, and a VLAN object with ID 14 under the first range, the VLAN object will belong to both VLAN range 1 and VLAN range 2. However, its parent VLAN range is the range under which it was created; in this case, the first VLAN range. The parent VLAN range is displayed in the **Parent VLAN Range** column when viewing a VLAN range details. If there are overlapping ranges in a VLAN view, you cannot disable the checkbox. If you want to disable the checkbox, you must reconfigure the VLAN ranges to avoid conflicts and then disable the option.
9. In the **Comment** field, enter comments if any.
10. Click **Next** to define extensible attributes. For more information, see [Managing Extensible Attributes](#).
11. Click **Save & Close** to save your configuration.

## Editing a VLAN View

To edit the details of a VLAN view:

1. From the **Data Management** tab, click the **VLANs** tab.

2. Select the VLAN view that you want to edit, select the Action icon and click **Edit**.  
**Note:** Changing the start or end VLAN ID values in a VLAN view may lead to errors. For example, if the VLAN IDs of a VLAN view are from 1 to 5, you cannot change them to 1 to 2 in case there is a child VLAN with ID 3 or 4 or 5 in the VLAN view.
3. Edit the details as required.
4. Click **Save & Close**.

## Deleting a VLAN View

You cannot delete a VLAN view that contains VLANs assigned to a network. To delete a VLAN view:

1. From the **Data Management** tab, click the **VLANS** tab.
2. Select the VLAN view that you want to delete, select the Action icon and click **Delete**.
3. Click **Yes** in the *Delete Confirmation* dialog box.



### Note

When you delete a VLAN view, all child VLANs and ranges will also be deleted.

## Configuring VLAN Ranges

A VLAN range is used by NIOS to group VLAN objects. A VLAN range ID must be in between 1 and 4094 (1 and 4094 included). A VLAN view can include one or more ranges. However, it is not mandatory that a VLAN view must include a range.

## Adding a VLAN Range

To add a VLAN range:

1. From the **Data Management** tab, click the **VLANS** tab.
2. From the Toolbar, click **Add-> VLAN Range**. The *Add VLAN Range Wizard* is displayed.
3. Click **Select VLAN View** to select the view to which the range must belong.
4. In the **VLAN Range Name** field, enter a name for the VLAN range.
5. In the **Start VLAN ID** field, enter the minimum ID that can be used by a VLAN object.
6. In the **End VLAN ID** field, enter the maximum ID that can be used by a VLAN object.
7. Select the **Pre-create VLANs** checkbox if you want NIOS to automatically create VLANs. When you select this checkbox, the **VLAN Name Prefix** field becomes enabled and NIOS creates VLANs between the range you specified in the **Start VLAN ID** and **End VLAN ID** fields. If you select this checkbox when adding a VLAN range, VLANs are created inside of the range. If you create an empty VLAN range, all existing VLANs (in a VLAN view) are re-parented by range if their IDs are in the start to end range IDs of the VLAN range. Therefore, if there are VLANs to re-parent by the newly created range, you cannot pre-create VLANs in this range. You cannot edit the **Pre-create VLANs** checkbox.
8. In the **VLAN Name Prefix** field, enter a prefix that will be added to the VLAN name. The VLAN name is constructed using the following pattern: `<VLAN Prefix Name> + <VLAN ID>` (the VLAN ID is optionally prefixed with zeros). VLANs are automatically created with the prefix name and VLAN ID. You cannot edit this field.
9. In the **Comment** field, enter additional information if any.
10. Click **Next** to define extensible attributes. For more information, see [Managing Extensible Attributes](#).
11. Click **Save & Close** to save your configuration.

## Editing a VLAN Range

To edit the details of a VLAN range:

1. From the **Data Management** tab, click the **VLANS** tab.

2. Click the VLAN view that the range belongs to, select the Action icon and click **Edit**.
3. Edit the details as required. You cannot edit the **VLAN Name Prefix** field and the **Pre-create VLANs** checkbox.
4. Click **Save & Close**.

**Note**

Changing the start and end VLAN IDs may lead to re-parenting of the child VLANs. Based on how the start and end VLAN IDs of the range changed, new VLANs may be added to the range or removed from it (in which case the VLAN view becomes the new parent). Also, changing parent field allows you to move ranges between views.

## Deleting a VLAN Range

To delete a VLAN range:

1. From the **Data Management** tab, click the **VLANs** tab.
2. Click the VLAN view that the range that you want to delete belongs to.
3. Select the VLAN range, select the Action icon and click **Delete**.
4. Select the **Delete only the VLAN Range and re-parent the VLANs** option if you want to delete only the VLAN range and not the VLAN child objects under the range. In this case, if you had selected the **Allow VLAN Range Overlapping** option when creating the VLAN view:
  - If the VLAN IDs of the VLAN objects can be categorized into another VLAN range, then the child objects are automatically categorized under the new VLAN range.
  - If the VLAN IDs cannot be categorized under a VLAN range, then the VLAN child objects are categorized under the VLAN view.
  - If the VLAN IDs can be categorized into more than one VLAN range, then NIOS displays an error message because it cannot determine the VLAN range under which to categorize the VLAN IDs and the VLAN range is not deleted.
5. Select the **Delete the VLAN Range and all VLANs below it** option if you want both the VLAN range and the VLAN child objects that belong to the VLAN range to be deleted.
6. Click **Yes** in the *Delete Confirmation* dialog box.

**Note**

You cannot delete or modify a VLAN object that is already assigned to a network. If a child VLAN object is assigned to a network and you select the Delete the VLAN Range and all VLANs below it option or the Delete only the VLAN Range and re-parent the VLANs option, NIOS displays an error message. You can resolve the error by manually restructuring the VLAN objects.

### Example - VLAN Range Deletion

The following example illustrates how VLAN objects are categorized when you delete a VLAN range.

```

VLAN View [1..100] (Allow Range Overlapping - ON)
|
--->VLAN Range 1 [20..80]
|   |
|   | ---> VLAN 30 (parent - Range 2)
|   |
|   | ---> VLAN 31
|   |
|   | ---> VLAN 40 (parent - Range 3)
|   |
|   | ---> VLAN 41 (parent - Range 3)
|
--->VLAN Range 2 [30..70]
|   |
|   | ---> VLAN 30
|   |
|   | ---> VLAN 31 (parent - Range 1)
|   |
|   | ---> VLAN 40 (parent - Range 3)
|   |
|   | ---> VLAN 41 (parent - Range 3)
|
--->VLAN Range 3 [40...60]
|
|   | ---> VLAN 40
|   |
|   | ---> VLAN 41

```

This example illustrates 3 VLAN ranges that allow overlapping VLAN IDs. If you try to delete VLAN Range 3, NIOS displays an error message because the child VLAN objects VLAN 40 and VLAN 41 can either be categorized under VLAN Range 1 or VLAN Range 2, and the VLAN Range 3 is not deleted. In this scenario, you must manually change the parent of the VLAN objects.

## Configuring VLAN Objects

You can create a VLAN topology or hierarchy by adding VLAN objects and grouping them together to create a VLAN view or range. You can then track statuses of the VLANs, set some additional details (such as department, contact, description). You can also assign VLANs to DHCP/IPAM networks.



## Viewing VLAN Objects in a VLAN View

To view all the VLAN objects that belong to a particular view:

1. From the **Data Management** tab, click the **VLANS** tab. A list of all the VLAN views is displayed. If the VLAN object belongs to a particular VLAN range, click the VLAN range.
2. Click the VLAN view whose details you want to see. All the VLAN objects and VLAN ranges belonging to the view are displayed with the following details:
  - **Name:** Name of the VLAN.
  - **Type:** VLAN or VLAN range.
  - **VLAN ID:** ID of the VLAN.
  - **Start VLAN ID:** Integer from which the VLAN view ID starts. This is applicable only to VLAN views and ranges.
  - **End VLAN ID:** Integer from which the VLAN view ID ends. This is applicable only to VLAN views and ranges.
  - **Description:** Description of the VLAN.
  - **Contact:** Contact information of the person or team managing or using the VLAN
  - **Department:** Name of the department using the VLAN
  - **Status:** Status can be **Assigned**, **Unassigned**, or **Reserved**. An Assigned status indicates that the VLAN is linked to an IPAM network object. An Unassigned status indicates that the VLAN is not linked to an IPAM object. A Reserved status indicates that the VLAN will be excluded from the list of VLANs that are proposed to be assigned to an IPAM object.
  - **Comment:** Any comment entered for the VLAN
  - **Assigned to:** IPAM objects that the VLAN is linked to. This is a read-only value and you cannot edit it.
  - Extensible attributes if any

### Note

The details you see vary depending on whether you are viewing the details of the VLAN view or the VLAN range.

Except the **Assigned to** value, you can double-click any of the VLAN details to edit them.

## Adding a VLAN Object

To add a VLAN object to your network:

1. From the **Data Management** tab, click the **VLANS** tab.
2. From the Toolbar, click **Add-> VLAN**.
3. In the *Add VLAN Wizard*, from the **VLAN Parent Type** drop-down list, select whether you want to add the VLAN to the VLAN view or VLAN range.
4. In the **VLAN Parent** field, select the specific VLAN view or range to which you want the VLAN to belong.
5. In the **VLAN Name** field, specify a name for the VLAN.
6. Click **Next Available VLAN ID** for NIOS to automatically assign an ID for the VLAN. You can also manually specify the VLAN ID. The value of the VLAN ID must be greater than or equal to 1 and lesser than or equal to 4094. When you click **Next Available VLAN ID**, NIOS searches for the lowest unused ID in the specified VLAN view or range.
7. Select the **Reserved** checkbox if you want the VLAN to be excluded from the list of VLANs that are proposed for assignment to an IPAM network object by clicking **Next Available VLAN**.
8. In the **Description** field, enter a description for the VLAN.
9. In the **Contact** field, enter the contact information of the person or team managing or using the VLAN. It can be a name, email ID, telephone number or so on.
10. In the **Department** field, enter the name of the department for which the VLAN will be used.
11. In the **Comment** field, enter additional information if any.
12. Click **Next** to define extensible attributes. For more information see [Managing Extensible Attributes](#).
13. Click **Save & Close** to save your configuration.

You can now add the VLAN to a VLAN view.

## Editing a VLAN Object

To edit the details of an existing VLAN object:

1. From the **Data Management** tab, click the **VLANs** tab.
2. From the list of VLAN views, click the VLAN view that the VLAN object to be edited belongs to. If the VLAN object belongs to a VLAN range, click the VLAN range.
3. Select the VLAN object that you want to edit, select the Action icon and click **Edit**.
4. Edit the details as required. You cannot edit the **Assigned to** field. You cannot change the ID, parent, and status for VLAN objects that are assigned to a network.
5. Click **Save & Close**.

## Editing Multiple VLAN Objects at the Same Time

You can select and edit multiple VLAN objects at the same time. To do this:

1. From the **Data Management** tab, click the **VLANs** tab.
2. From the list of VLAN views, click the VLAN view that the VLAN objects to be edited belong to. If the VLAN objects belong to a VLAN range, click the VLAN range.
3. Select all the VLAN objects that you want to edit, select the Action icon and click **Edit**.
4. Edit the details as required. You can edit the parent VLAN, **Reserved** checkbox, **Description**, **Contact**, **Department**, and **Comment** fields. However, you cannot edit the parent of a VLAN object that does not belong to a VLAN range and belongs directly to a VLAN view.
5. Click **Save & Close**.

### Notes

- You can select multiple VLAN objects to edit only if all the selected VLAN objects belong to the same VLAN range.
- You can edit the **VLAN Parent** field to select only another VLAN range as the new parent. You cannot select a VLAN view as the new parent. The new VLAN range must belong to the same VLAN view.
- You cannot edit the **VLAN Parent** field for multiple VLAN objects that belong directly to a VLAN view and not to a VLAN range.
- You cannot edit the VLAN parent of VLAN objects that are assigned to a network.

## Deleting a VLAN Object

You cannot delete a VLAN object that is assigned to a network. For more information, see [Assigning VLANs to a Network](#).

To delete a VLAN object:

1. From the **Data Management** tab, click the **VLANs** tab.
2. From the list of VLAN views, click the VLAN view that the VLAN to be deleted belongs to. If the VLAN object belongs to a VLAN range, click the VLAN range.
3. Select the VLAN object that you want to delete, select the Action icon and click **Delete**.
4. Click **Yes** in the *Delete Confirmation* dialog box.

## Assigning VLANs to a Network

You can assign a VLAN object to different networks. You can also assign multiple VLAN objects to a single network. You can also assign multiple networks to a single VLAN. Once a VLAN object is assigned to a network, the status of VLAN

object changes from Unassigned or Reserved to **Assigned**. For information about the types of VLAN status, see [Configuring VLANs](#).

Assigning a VLAN object enables you to compare an assigned VLAN object with the VLAN objects discovered by Network Insight. You can then generate inventory and conflict reports based on this data.

You can also create [smart folders](#) and obtain a list of networks by the assigned VLAN ID or name.



#### Note

- You can assign VLAN objects only to IPv4 or IPv6 networks.
- You cannot assign VLAN objects to a network container.
- If a VLAN object is assigned to a network, you cannot convert that network into a network container.
- You cannot split or join a network that has VLAN objects assigned to it.

## Viewing VLAN Objects Assigned to a Network

You can view the VLAN objects assigned to a network on the **IPAM** tab or the **DHCP** tab. The **Assigned VLAN ID** and **Assigned VLAN Name** columns display the VLAN ID and name of the assigned object. If multiple VLANs are assigned to a network, **Multiple** is displayed in the **Assigned VLAN ID** and **Assigned VLAN Name** columns. You can then compare it to the data in the **Discovered VLAN ID** and **Discovered VLAN Name** columns and view conflicts if any.

## Adding VLAN Objects to a Network

To add or assign a VLAN object to a network, perform the following steps:

1. From the **Data Management** tab, select the **IPAM** tab or **DHCP** tab -> *network* checkbox of the network to which you want to assign the VLAN object.
2. Select the Action icon and then click **Edit**.
3. Click **VLAN Assignment**. A list of all the VLAN objects associated with the network is displayed.
4. Click the Add icon.
5. From the **VLAN Parent Type** drop-down list, select the parent type of the VLAN object that you want to add. The parent type can either be a VLAN view or a VLAN range. For more information, see [Configuring VLANs](#).
6. From the **VLAN Parent** field, select the VLAN view or VLAN range that the VLAN object belongs to.
7. From the **VLAN Name** field, select the VLAN object. Alternatively, click **Next Available VLAN** if you want NIOS to assign a VLAN object. NIOS searches for a VLAN object with the lowest ID in the specified VLAN view or range. VLAN objects whose status is **Assigned** or **Reserved** are ignored. However, you can still select a reserved or assigned VLAN object by clicking the **VLAN Name** field and selecting the VLAN object. The **VLAN ID** field is populated with the ID of the selected VLAN object.
8. Click **Add** to add the VLAN objects. The table is populated with the VLAN object details.
9. You can click the Add icon to add more VLAN objects.
10. Click **Save & Close**.

You can also assign a VLAN object when creating an IPv4 or IPv6 network. For more information, see [Configuring IPv4 Networks](#) and [Managing IPv6 Networks](#).

## Editing Assigned VLAN Objects

To edit a VLAN object that is already assigned, perform the following steps:

1. From the **Data Management** tab, select the **IPAM** tab or **DHCP** tab -> *network* checkbox of the network that you want to edit.
2. Select the Action icon and then click **Edit**.
3. Click the **VLAN Assignment** tab.
4. Select the VLAN object that you want to edit and click the Edit icon. Details of the VLAN object are displayed. You can edit and select another VLAN object in another VLAN view for example.
5. Click **Save & Close**.

## Deleting Assigned VLAN Objects

You cannot delete a VLAN object that is assigned to a network. You must first unassign the VLAN object and then delete it. Deleting an unassigned VLAN object from the list of VLAN objects that are assigned to a network does not delete it from NIOS. Also, you cannot delete a network, network view, or a parent network container that have VLAN objects assigned to them.

To unassign a VLAN object, perform the following steps:

1. From the **Data Management** tab, select the **IPAM** tab or **DHCP** tab -> *network* checkbox of the network that you want to edit.
2. Select the Action icon and then click **Delete**.
3. Click the **VLAN Assignment** tab.
4. Select the VLAN object that you want to edit and click the Delete icon. The VLAN object is deleted from the assignment.
5. Click **Save & Close**.

## Using the NIOS CLI

The Infoblox NIOS CLI (Command Line Interface) allows you to configure and monitor the appliance from a remote console using a set of commands. Some administrative tasks, such as resetting the appliance, can be done only through the CLI. CLI commands do not support IDNs. These commands display IDN data in punycode only. For more information about IDN, see [Multilingual Support](#).

This section explains the CLI commands that you can use to configure and manage the NIOS appliance from a remote terminal. For the latest Infoblox documentation, visit the Infoblox Support web site at <https://support.infoblox.com/>.

Topic	Content
Overview	Explains how to access the Infoblox CLI using a console port or SSHv2 client. This topic also describes the CLI conventions and outlines the basic CLI commands.
CLI Commands	Explains the function and usage of each command, and provides an example of the command usage and expected results.

## Conventions

The conventions used in this section follow the Infoblox documentation style conventions, as listed in the following table.

Style	Usage
<code>screen</code>	Indicates session text or system information displayed on the screen
<b>boldface screen</b>	Signifies command line entries that you type.
<i>italic screen</i>	Signifies variables that you enter for your configuration, such as file names and group names.

CLI syntax uses conventions that are unique to documenting command line tools. The following table provides a list of syntax delimiters and their meanings.

Item	Convention
{ } brackets	Indicates a mandatory feature.
[ ] brackets	Indicates an optional feature.
pipe symbol	Indicates an “or” relationship between two features.

## Variables

Infoblox uses the following variables to represent the values of the configurations that exist on your appliance. You should substitute the variables with the actual values that match your site configuration.

Variable	Value
admin_group	Name of a group of administrators
admin_name	Name of the appliance administrator
addr_range	IP address range
domain_name	Domain name
directory	Directory name
dns_view	DNS view
filter_name	Filter name
grid_master	Grid master
grid_member	Grid member
hostname	Host name of an independent appliance
id_grid	Grid name
ip_addr	IPv4 address
member	Grid member name
netmask	Subnet mask
network	IP address of a network
numerical	Numerical entry
zone	DNS zone

The following is a list of commands that NIOS supports:

- [Accessing the Infoblox CLI](#)
- [CLI Commands](#)
- [?](#)
- [ddns\\_add](#)
- [ddns\\_delete](#)
- [delete](#)
- [dig](#)
- [dns\\_a\\_record\\_delete](#)
- [exit](#)
- [help](#)

- *ping*
- *quit*
- *reboot*
- *reset all*
- *reset arp*
- *reset cli*
- *reset database*
- *reset reporting\_data*
- *reset snmp*
- *reset ssh\_keys*
- *restart service*
- *rotate log*
- *set admin\_group\_acl*
- *set adp*
- *set allow\_query\_domain*
- *set apache\_https\_cert*
- *set auto\_provision*
- *set bfd*
- *set bgp log*
- *set bloxtools*
- *set cc\_mode*
- *set certificate\_auth\_admins*
- *set certificate\_auth\_services*
- *set check\_auth\_ns*
- *set circ\_txn\_trace*
- *set cloud\_services\_portal\_force\_refresh*
- *set connection\_limit*
- *set debug*
- *set debug\_analytics*
- *set default\_revert\_window*
- *set default\_route*
- *set delete\_tasks\_interval*
- *set dhcpd\_recv\_sock\_buf\_size*
- *set extra\_dns\_name\_validations*
- *set disable\_gui\_one\_click\_support*
- *set disable\_https\_cert\_regeneration*
- *set dns*
- *set dns-accel*
- *set dns\_rrl*
- *set docker\_bridge*
- *set dscp*
- *set enable\_dnstap*
- *set enable\_match\_recursive\_only*
- *set fips\_mode*
- *set hardware-type*
- *set ibtrap*
- *set interface*
- *set ip\_rate\_limit*
- *set ipam\_web\_ui*
- *set ipv6\_disable\_on\_dad*
- *set ipv6\_neighbor*
- *set ipv6\_ospf*
- *set ipv6\_status*

- *set lcd keys or set lcd*
- *set lcd\_settings*
- *set lcd\_settings hwident*
- *set license*
- *set lines*
- *set log\_txn\_id*
- *set lom*
- *set max\_recursion\_depth*
- *set max\_recursion\_queries*
- *set membership*
- *set mgm attached*
- *set mld\_version\_1*
- *set monitor dns*
- *set monitor dns alert*
- *set ms\_dns\_reports\_sync\_interval*
- *set ms\_sticky\_ip*
- *set named\_recv\_sock\_buf\_size*
- *set named\_tcp\_clients\_limit*
- *set network*
- *set nogrid*
- *set nomastergrid*
- *set nosafemode*
- *set oosp*
- *set ospf*
- *set overload\_bootp*
- *set pc\_domain add*
- *set pc\_domain delete*
- *set phonehome*
- *set promote\_master*
- *set gmc\_promotion*
- *set prompt*
- *set regenerate\_anycast\_password*
- *set remote\_console*
- *set reporting\_cert*
- *set reporting\_user\_capabilities*
- *set restart\_anycast\_with\_dns\_restart*
- *set revert\_grid*
- *set rpz\_recursive\_only*
- *set safemode*
- *set scheduled*
- *set security*
- *set session\_timeout*
- *set smartnic monitor-mode*
- *set snmptrap*
- *set static\_route*
- *set subscriber\_secure\_data add*
- *set subscriber\_secure\_data bypass*
- *set subscriber\_secure\_data delete*
- *set subscriber\_secure\_data clear\_all*
- *set subscriber\_secure\_data garbage\_collect*
- *set subscriber\_secure\_data never\_proxy*
- *set subscriber\_secure\_data persist*
- *set support\_access*



- *set support\_timeout*
- *set sysName*
- *set tcp\_timestamps*
- *set temp\_license*
- *set term*
- *set test\_promote\_master*
- *set thresholdtrap*
- *set token*
- *set traffic\_capture*
- *set txn\_trace*
- *set update\_rabbitmq\_password*
- *set upgrade\_dist\_rsync\_batch disable*
- *set wins\_forwarding*
- *show admin\_group\_acl*
- *show allow\_query\_domain*
- *show allow\_query\_domain\_views*
- *show analytics\_parameter*
- *show adp*
- *show arp*
- *show action\_to\_activate\_hotfix*
- *show auto\_provision*
- *show bfd*
- *show bgp*
- *show bloxtools*
- *show capacity*
- *show cc\_mode*
- *show certificate\_auth\_admins*
- *show certificate\_auth\_services*
- *show check\_auth\_ns*
- *show clusterd\_info*
- *show config*
- *show connections*
- *show connection\_limit*
- *show cpu*
- *show date*
- *show debug*
- *show debug\_analytics*
- *show delete\_tasks\_interval*
- *show dhcp\_gss\_tsig*
- *show dhcpd\_recv\_sock\_buf\_size*
- *show dhcpv6\_gss\_tsig*
- *show disk*
- *show dns*
- *show dns-accel*
- *show dns-accel-cache*
- *show dns\_gss\_tsig*
- *show dns-over-tls-config*
- *show dns-over-tls-stats*
- *show dns-over-tls-status*
- *show dns\_rrl*
- *show dnstap-stats*
- *show dnstap-status*
- *show docker\_bridge*

- *show doh-config*
- *show doh-stats*
- *show doh-status*
- *show dscp*
- *show dtc\_geoip*
- *show enable\_match\_recursive\_only*
- *show extra\_dns\_name\_validations*
- *show file*
- *show fips\_mode*
- *show hardware\_status*
- *show hardware-type*
- *show hwid*
- *show ibtrap*
- *show interface*
- *show ip\_rate\_limit*
- *show ipv6\_bgp*
- *show ipv6\_disable\_on\_dad*
- *show ipv6\_neighbor*
- *show ipv6\_ospf*
- *show lcd*
- *show lcd\_info*
- *show lcd\_settings*
- *show license*
- *show license\_uid*
- *show license\_pool\_container*
- *show log\_guest\_lookups*
- *show log\_txn\_id*
- *show lom*
- *show max\_recursion\_depth*
- *show max\_recursion\_queries*
- *show memory*
- *show mld\_version*
- *show monitor*
- *show monitor dns alert*
- *show monitor dns alert status*
- *show ms\_sticky\_ip*
- *show named\_rcv\_sock\_buf\_size*
- *show network*
- *show ntp*
- *show ospf*
- *show overload\_bootp*
- *show pc\_domain*
- *show phonehome*
- *show query\_capture*
- *show remote\_console*
- *show reporting\_user\_capabilities*
- *show restart\_anycast\_with\_dns\_restart*
- *show routes*
- *show rpz\_recursive\_only*
- *show scheduled*
- *show security*
- *show session\_timeout*
- *show smartnic*

- *show snmp*
- *show static\_routes*
- *show status*
- *show subscriber\_secure\_data*
- *show subscriber\_secure\_data bypass*
- *show subscriber\_secure\_data garbage\_collect*
- *show subscriber\_secure\_data cache\_usage\_stats*
- *show subscriber\_secure\_data never\_proxy*
- *show subscriber\_secure\_data persist*
- *show support\_access*
- *show support\_timeout*
- *show tcp\_timestamps*
- *show tech-support*
- *show temperature*
- *show test\_promote\_master*
- *show thresholdtrap*
- *show token*
- *show traffic\_capture\_status*
- *show upgrade\_history*
- *show uptime*
- *show version*
- *show vpn\_cert\_dates*
- *show wins\_forwarding*
- *shutdown*
- *snmpget*
- *snmpwalk*
- *traceroute*
- *show rabbitmq\_queues*

## Accessing the Infoblox CLI

You can access the Infoblox CLI from a management system. The management system is the computer from which you configure and monitor the NIOS appliance. You can access the Infoblox CLI from the management system directly through a serial cable or remotely across an Ethernet network.

- Console port access—Access the Infoblox CLI through a direct console connection from your management system to the appliance.
- SSHv2 client access—Accessing the Infoblox CLI remotely by making an SSHv2 connection across an Ethernet network.

### Note

Only superusers can log in to the appliance through a console connection.

### Console Port Access

You can access the Infoblox CLI by using a terminal emulation program from the management system through a direct console connection.

To access the Infoblox CLI through the console port:

1. Connect a serial cable from the console port on your management system to the console port on the appliance. The appliance has a male DB-9 console port on its front panel.
2. Use the following connection settings to launch an emulation session through a serial terminal emulation program such as Hilgraeve Hyperterminal® (provided with the Windows® operating systems):

- Bits per second: 9600
  - Data bits: 8
  - Parity: None
  - Stop bits: 1
  - Flow control: Xon/Xoff
3. Use the following default user name and password to log in to the Infoblox appliance:

**admin**

infoblox

#### Note

User names and passwords are case-sensitive.

### SSHv2 Client Access

You can access the Infoblox CLI from a remote management system. You must first enable remote console access before you can remotely access the Infoblox CLI. By default, remote console access (SSHv2 access) is disabled on the Infoblox appliance.

You can enable remote console access on the Infoblox appliance through either the Infoblox GUI or the CLI. To enable remote console access through the Infoblox GUI:

1. Make an HTTPS or console connection to the appliance and log in to the appliance.
2. For a Grid member or Grid Master, complete the following:
  1. From the **Grid** tab, select the **Grid Manager** tab, and then click **Grid Properties** -> **Edit** from the Toolbar.
  2. In the *Grid Properties* editor, select the **Security** tab, and then select **Enable Remote Console Access**.
  3. Click **Save & Close**.

For an independent appliance, complete the following:

1. From the **System** tab, select the **System Manager** tab, and then click **System Properties** -> **Edit** from the Toolbar.
2. In the *System Properties* editor, select the **Security** tab, and then select **Enable Remote Console Access**.
3. Click **Save&Close**.

To enable remote console access through the CLI:

1. From the command line, enter the following after the *Infoblox >* prompt:
 

```
set remote_console
```
2. Enter **y** at the *Enable remote console access (grid-level)? (y or n):* prompt.
3. Confirm the settings.

After you enable the remote console access, you can access the Infoblox CLI from a remote location using an SSHv2 client.

To access the Infoblox CLI using an SSHv2 client:

1. On the management system, open a remote console connection through an SSHv2 client.
2. In a shell window or terminal window, log in with an account that has superuser privileges.
3. Enter the user name and host name or IP address of the appliance. For example:
 

```
ssh admin@192.168.1.2
```
4. Optionally, you can launch a graphical SSHv2 client and enter the information in the appropriate fields.

## CLI Commands

The basic Infoblox CLI commands are alphabetically listed in the following table.

CLI Command	Description
?	Displays the help information.
ddns_add	Sends DDNS updates to add records.
ddns_delete	Sends DDNS updates to delete records.
delete	Deletes specific files.
dig	Performs a DNS lookup and prints the results.
dns_a_record_delete	Delete a DNS A record
exit	Exits the command interpreter.
help	Displays the help information.
ping	Sends ICMP ECHO requests to verify that the host is functioning properly.
quit	Exits the command interpreter.
reboot	Reboots the Infoblox appliance.
reset	Resets the system settings.
rotate	Rotates specific files.
set	Sets the current system settings. This command has other related commands.
show	Shows the current system settings. This command has other related commands.
shutdown	Shuts down the Infoblox appliance.
traceroute	Displays the path or route diagnostic information of the IPv4/IPv6 packets.

The

**reset**

, **set**, and **show** commands each have related commands. To view a complete list of the related commands on the remote console, go to the command prompt and enter

**help set**

or

**help show**

.

**Note**

For the CLI `set traffic_capture` password field will support only these special characters, A-Z, a-z, 0-9, [, ], ,, \_(underscore), /, \, #, &, @, %, ,(Comma), ~, !, \$, `(grave accent), {, }, (, ), =, - (dash), \*, ^, +, ?.

## Using CLI Help

You can display a list of available CLI commands by typing `help` at the command prompt. For example:

```
> help
?          Display help
ddns_add   Send DDNS update to add a record
ddns_delete Send DDNS update to delete a record
delete     Delete files
dig        Perform a DNS lookup and print the results
dns_a_record_delete Delete a DNS A record
exit       Exit command interpreter
help       Display help
ping       Send ICMP ECHO
quit       Exit command interpreter
reboot     Reboot device
reset      Reset system settings
rotate     Rotate files
set        Set current system settings
show       Show current system settings
shutdown   Shutdown device
traceroute Route path diagnostic
```

To view a detailed explanation about a CLI command and its syntax, type `help <command>` after the command prompt. For example:

```
> help rotate
```

**Synopsis:**

```
rotate log [ syslog | debug | audit | ifmapserver]
rotate file groupname filename [ filename2, filename3, ...]
```

**Description:**

Rotates the specified log file, up to 10 previous.  
logfiles will be preserved

## ?

The ? command displays information about a specified CLI command. If you do not specify a command, a list of all available commands is shown.

## Syntax

? [command]

Argument	Description
<code>command</code>	A variable that you substitute with any CLI command to display a description of the function and a synopsis of its usage.

## Examples

### Display a list of commands

```
Infoblox > ?
=====
Command Summary
=====
?                Display help
deleteDelete    files
dig             Perform a DNS lookup and print the results
exit           Exit command interpreter
help           Display help
ping           Send ICMP ECHO
quit          Exit command interpreter
reboot        Reboot device
reset         Reset system settings
set           Set current system settings
show         Show current system settings
shutdown     Shutdown device
traceroute   Route path diagnostic
ddns_add     Send DDNS update to add a record
```

```
ddns_delete          Send DDNS update to delete a record
rotate              Rotate files
```

```
=====
```

## Display details for a single command

```
Infoblox > ? exit
```

Synopsis:

```
exit, quit
```

Description:

```
Exits the command interpreter. There are no arguments to exit.
```

## ddns\_add

The **ddns\_add** command sends DDNS updates to the appliance when you add new resource records. To use this command, ensure that you have properly configured the appliance for DDNS updates. For information, refer to the *Infoblox NIOS Administrator Guide*. To update a record that contains IDN, you must enter the domain name in punycode. The appliance retains the record in punycode and does not convert punycode to IDN.

## Syntax

```
ddns_add <domain-name> <ttl> <type> <data> [keyname:secret]
```

Argument	Description
domain-name	The FQDN of the resource record being added. For example, if the name of the record is <b>dns1</b> and the forward-mapping zone name is <b>corp100.com</b> , the FQDN is <b>dns1.corp100.com</b> . For an IDN, use the punycode representation of the IDN. For example, if the name of the record is 域 and the forward-mapping zone name is <b>corp100.com</b> , the FQDN is <b>xn--cjs.corp100.com</b> .
ttl	The TTL value (in seconds) of the new resource record.
type	The record type of the new resource record. For example, enter A for an A record and PTR for a PTR record.
data	The RDATA of the resource record. For an IDN, use the punycode representation of the IDN. For example, enter the IP address of an A record or the domain name of a PTR record.
[keyname:secret]	The TSIG key name and the secret for sending DDNS updates. You must enter the TSIG key name and shared secret if the DNS zone to which the record belongs is configured with a TSIG key.



## Example

```
Infoblox > ddns_add dns1.corp100.com 20 A 10.0.0.11
```

## ddns\_delete

The **ddns\_delete** command sends DDNS updates to the appliance when you delete existing resource records. To use this command, ensure that you have properly configured the appliance for DDNS updates. For information, refer to the *Infoblox NIOS Administrator Guide*. To delete a record that contains IDN, you must enter the domain name in punycode.

## Syntax

```
ddns_delete <domain-name> [type[keyname:secret]]
```

Argument	Description
<code>domain-name</code>	The FQDN of the resource record being deleted. For example, if the name of the record is <code>dns1</code> and the forward-mapping zone name is <code>corp100.com</code> , the FQDN is <code>dns1.corp100.com</code> .  For an IDN, use the punycode representation of the IDN. For example, if the name of the record is 域 and the forward-mapping zone name is <code>corp100.com</code> , the FQDN is <code>xn--cjs.corp100.com</code> .
<code>type</code>	The record type of the resource record. For example, enter <code>A</code> for an <code>A</code> record and <code>PTR</code> for a <code>PTR</code> record. This is optional.
<code>[keyname:secret]</code>	The TSIG key name and the secret for sending DDNS updates. You must enter the TSIG key name and shared secret if the DNS zone to which the record belongs is configured with a TSIG key.

## Example

```
Infoblox > ddns_delete dns1.corp100.com
```

## delete

The **delete** command deletes a specific file or a set of files.

## Syntax

```
delete [file]
```

Argument	Description
<code>file</code>	The name of the file which needs to be deleted

## Example

Infoblox > **delete abc.csv**

## dig

The **dig** command performs a DNS lookup on a specified server and displays the results. You can also use the **inverse** command to perform a reverse DNS lookup. This command displays IDN data in punycode, if any, for the specified server. If you specify IP address of the Microsoft server in this command, the IDN data is displayed in \xyz format.

## Syntax

```
dig [@server_address] <hostname> [type] [opt...]
```

```
dig [@server_address] <ip-address> inverse
```

Argument	Description
server_address	The IP address of the host on which you want to perform a DNS lookup.
hostname	The name of the host on which you want to perform a DNS lookup.
ip-address	The IP address of the host on which you want to perform a DNS lookup.
type	You can enter any of the following for the object type (case sensitive): a, a6, aaaa, afsdb, any, apl, axfr, cert, cname, dhcid, div, dname, dnskey, ds, gpos, hinfo, hip, ipseckey, isdn, ixfr, key, keydata, kx, loc, maila, mailb, mb, md, mf, mg, minfo, mr, mx, naptr, none, ns, nsap, nsap_ptr, nsec, nsec3, nsec3param, null, nxt, opt, ptr, px, rp, rrsig, rt, sig, soa, spf, srv, sshfp, tkey, tsig, txt, unsec, wks, and x25. The default is a.

Argument	Description
opt	<p>You can enter one or more of the following options:</p> <ul style="list-style-type: none"> <li>• <code>-x</code> (specifies the in-addr lookup)</li> <li>• <code>-b address</code> (specifies the binding to the source address)</li> <li>• <code>-y name:key</code> (specifies the named base64 tsig key)</li> <li>• <code>+vc</code> (enables the TCP mode)</li> <li>• <code>+norecurse</code> (disables the recursive mode)</li> <li>• <code>+short</code> (disables everything except the short forms of answers)</li> <li>• <code>+nssearch</code> (searches all the authoritative nameservers)</li> <li>• <code>+trace</code> (traces all the delegations from the root)</li> <li>• <code>+cdflag</code> (requests the server not to perform a DNSSEC validation)</li> <li>• <code>+dnssec</code> (requests the server to send DNSSEC records)</li> <li>• <code>+multiline</code> (displays records in multiple lines)</li> </ul>

## Examples

### Perform a DNS lookup

```

Infoblox > dig @10.0.2.60 www.infoblox.com a
: <<>> DiG 9.6.1-p3 <<>> @10.0.2.60 -x www.infoblox.com a
: <1 server found>
:: global options: +cmd
:: Got answer:
:: >>HEADER<< opcode: QUERY, status: NOERROR, id: 45283
:: flags: qr aa rd ra: QUERY: 1, ANSWER: 2, AUTHORITY: 0, ADDITIONAL: 0
:: QUESTION SECTION:
: www.infoblox.com.                IN      A
:: ANSWER SECTION:
www.infoblox.com      3600    IN      CNAME   infoblox.com.
infoblox.com          600     IN      A
128.242.99.236
:: Query time: 2 msec
:: SERVER: 10.0.2.60#53<10.0.2.60>
:: WHEN: Fri Feb 26 14:06:00 2010
:: MSG SIZE rcvd: 64

```

## Perform a reverse DNS lookup

```
Infoblox > dig @10.0.2.60 inverse
: <<>> DiG 9.6.1-p3 <<>> @10.0.2.60 inverse
: <1 server found>
:: global options: +cmd
:: Got answer:
:: >>HEADER<< opcode: QUERY status: NXDOMAIN, id: 37916
:: flags: qr rd ra: QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADDITIONAL: 0
:: QUESTION SECTION:
:inverse.                IN          A
:: AUTHORITY SECTION:
. 10800                  IN          SOA        a.root-servers.net.
nstld.verisign-grs.com. 2010022601 1800 900 604800 86400
:: Query time: 132 msec
:: SERVER: 10.0.2.60#53<10.0.2.60>
:: WHEN: Thu Feb 25 11:20:09 2010
:: MSG SIZE rcvd: 100
```

## dns\_a\_record\_delete

The **dns\_a\_record\_delete** command deletes a DNS record of A record type.

## Syntax

```
dns_a_record_delete <dns view> <record name> <zone name> <ip address>
```

Argument	Description
<code>&lt;dns view&gt;</code>	The DNS view where the target DNS A record belongs.
<code>&lt;record name&gt;</code>	The name of the target DNS A record.
<code>&lt;zone name&gt;</code>	The name of the parent zone.
<code>&lt;ip address&gt;</code>	The IP address of the target record.

## Example

```
Infoblox > dns_a_record_delete default_view my_record.with.long.name test.com 1.2.3.4
```

## exit

The **exit** (**quit**) command terminates the command line interface and halts the CLI session.

## Syntax

```
exit, quit
```

Both commands produce the same results. There are no arguments for either command.

Command	Description
exit	Terminates the current CLI session.
quit	Terminates the current CLI session.

## Examples

```
Infoblox > exit
Good Bye
Connection to <IP address> closed.
Infoblox > quit
Good Bye
Connection to <IP address> closed.
```

## help

The **help** command displays information about a specified CLI command. If you do not specify a command, a list of all available commands is shown.

## Syntax

```
help [command]
```

Argument	Description
<code>command</code>	A variable that you substitute with any CLI command to display a description of the function and a synopsis of its usage.

## Examples

### Display a list of commands

```
Infoblox > help
=====
Command Summary
=====
?                Display help
deleteDelete    files
dig              Perform a DNS lookup and print the results
exit            Exit command interpreter
help            Display help
ping            Send ICMP ECHO
quit            Exit command interpreter
reboot          Reboot device
reset           Reset system settings
set             Set current system settings
show           Show current system settings
shutdown        Shutdown device
traceroute      Route path diagnostic
ddns_add        Send DDNS update to add a record
ddns_delete     Send DDNS update to delete a record
rotate          Rotate files
=====
```

### Display details for a single command

```
Infoblox > help exit
Synopsis:
  exit, quit

Description:
  Exits the command interpreter. There are no arguments to exit.
```

# ping

The **ping** command verifies if a remote IPv4/IPv6 host is functioning and accessible across the network. When you execute the ping command, it sends five (default) sequential ICMP ECHO requests to the host and displays the results.

## Syntax

```
ping {hostname | ip_address} [ opt ]
```

Argument	Description
hostname	The name of the remote host that you want to verify.
ip_address	The IP address of the remote host that you want to verify.
opt	<ul style="list-style-type: none"><li>• <b>numerical</b> (specifies to not interpret the IP address as a DNS name)</li><li>• <b>src_addr</b> (specifies the starting or "from" address)</li><li>• <b>v6</b> (specifies you are using an IPv6 hostname)</li><li>• <b>broadcast</b> (allows pinging to a broadcast address)</li><li>• <b>ttl&lt;hops&gt;</b> (specifies the time-to-live setting for outgoing packets)</li><li>• <b>packetsize&lt;bytes&gt;</b> (specifies the number of data bytes to send)</li><li>• <b>count&lt;packets&gt;</b> (specifies number of echo_requests packets sent, default is 5, maximum is 250)</li></ul>

## Examples

### Valid host

```
Infoblox > ping 10.1.1.1
pinging 10.1.1.1
PING 10.1.1.1 (10.1.1.1) 56(84) bytes of data.
64 bytes from 10.1.1.1: icmp_seq=1 ttl=64 time=0.295 ms
64 bytes from 10.1.1.1: icmp_seq=2 ttl=64 time=0.102 ms
64 bytes from 10.1.1.1: icmp_seq=3 ttl=64 time=0.155 ms
64 bytes from 10.1.1.1: icmp_seq=4 ttl=64 time=0.211 ms
64 bytes from 10.1.1.1: icmp_seq=5 ttl=64 time=0.265 ms
```

```
- 10.1.1.1 ping statistics -  
5 packets transmitted, 5 received, 0% packet loss, time 4005ms rtt min/avg/  
max/mdev = 0.335/0.562/1.245/0.343 ms
```

## Invalid host

```
Infoblox > ping jsparrow  
pinging jsparrow  
ping: unknown host jsparrow
```

## quit

The **quit** command terminates the command line interface and halts the CLI session.

## Syntax

```
quit, exit
```

Both commands produce the same results. There are no arguments for either command.

Command	Description
quit	Terminates the current CLI session.
exit	Terminates the current CLI session.

## Examples

```
Infoblox > quit  
Good Bye  
Connection to <IP address> closed.
```

```
Infoblox > exit  
Good Bye  
Connection to <IP address> closed.
```

## reboot

The **reboot** command halts and then restarts the appliance. Use this command as a last measure when the appliance appears to be hung. Rebooting the appliance clears the cache and resets the system.



## Syntax

```
reboot
```

There are no arguments for this command.

## Example

```
Infoblox > reboot  
REBOOT THE SYSTEM? (y or n) y
```

## reset all

The **reset all** command clears the NIOS appliance of database, configuration, and network settings. It then re-establishes the factory settings with the default IP address, gateway, and subnet mask.

The **reset all licenses** command clears database, configuration, and network settings. It also clears all licensing information from the appliance before re-establishing the factory settings.

The **reset all auto\_provision** command clears database, configuration, and network settings. It also re-enables auto-provisioning and a dynamic IP address is assigned to the appliance.



### Note

No previous data remains on the appliance after using these commands.

## Syntax

```
reset all [licenses | auto_provision]
```

Argument	Description
licenses	Specifies the removal of all licenses during the process of re-establishing the factory settings on the appliance.
auto_provision	Auto-provisioning is re-enabled and a dynamic IP address is assigned after re-establishing the factory settings on the appliance.

## Examples

### Re-establish factory settings

```
Infoblox > reset all  
The entire system will be reset to default settings.
```

```
WARNING: THIS WILL ERASE ALL DATA AND LOG FILES THAT HAVE BEEN CREATED ON THIS SYSTEM. ARE YOU SURE YOU WANT TO PROCEED? (y or n): y
```

## Re-establish factory settings and remove all licenses

```
Infoblox > reset all licenses
```

```
The entire system will be reset to default settings and all licenses will be removed. WARNING: THIS WILL ERASE ALL DATA AND LOG FILES THAT HAVE BEEN CREATED ON THIS SYSTEM. ARE YOU SURE YOU WANT TO PROCEED? (y or n): y
```

## Re-establish factory settings and re-enable auto-provisioning

```
Infoblox > reset all auto_provision
```

```
The entire system will be reset to default settings and system will try to obtain a dynamic address.
```

```
WARNING: THIS WILL ERASE ALL DATA AND LOG FILES THAT HAVE BEEN CREATED ON THIS SYSTEM. ARE YOU SURE YOU WANT TO PROCEED? (y or n): y
```

## reset arp

The **reset arp** command clears the ARP (Address Resolution Protocol) cache. The ARP maps IP addresses to the hardware MAC addresses and logs them in a table which is stored in the cache. Over time, the IP address leases expire and are assigned to new devices (MAC addresses). Infoblox recommends that you periodically clear this cache to maintain valid mappings between IP addresses and MAC addresses.

## Syntax

```
reset arp
```

This command has no arguments.

## Example

```
Infoblox > reset arp
```

```
ARP cache cleared.
```

## reset cli

This command clears obsolete CLI credentials (community strings) of devices polled by Network Insight. After that, Network Insight reguesses the CLI credentials for each device.

You can run this command on the following Grid members:

- Grid Master
- Discovery Consolidator

- Discovery Probes

When the command is executed on the Grid Master or discovery Consolidator, the CLI credentials are reset for all devices on all discovery Probes. When the command is executed on a discovery Probe, the CLI credentials are reset for devices only on this Probe.

## Syntax

```
reset cli
```

This command has no arguments.

## Example

```
Infoblox > reset cli
```

```
Resetting CLI credentials will clear all assigned CLI credentials and cause to
reguess them for each device.
```

```
Do you wish to continue to the CLI credentials (y or n): y
```

```
Probe probe.blox: Success
```

```
The command 'reset cli' completed successfully.
```

```
Infoblox >
```

## reset database

The **reset database** command removes configuration files and DNS and DHCP data from the NIOS appliance database. However, the network and licensing information remains intact. The network settings of the appliance include the IP address and subnet mask for the appliance, the IP address of the gateway, the host name, and the remote access setting.

You can use this command to diagnose problems such as the following:

- Misplacing the administrator account and password.
- Preserving the log files when clearing the database.

The **reset database auto\_provision** command resets the NIOS appliance to default settings, re-enables auto-provisioning, and a dynamic IP address is assigned to the appliance.

## Syntax

```
reset database [auto_provision]
```

Argument	Description
auto_provision	Auto-provisioning is re-enabled and a dynamic IP address is assigned after resetting the database of appliance.

## Reset the database

```
Infoblox > reset database  
The following network settings can be restored after reset: IP Address:  
10.1.1.10  
Subnet Mask: 255.255.255.0  
Gateway: 10.1.1.1  
Host Name: ns1.corp100.com  
Remote Console Access: true The entire database will be erased. Do you wish  
to preserve basic network settings? (y or n) y
```

## Reset the database and re-enable auto-provisioning

```
Infoblox > reset database auto_provision  
The entire system will be reset to default settings and system will try to  
obtain a dynamic address.  
WARNING: THIS WILL ERASE ALL DATA AND LOG FILES THAT HAVE BEEN CREATED ON  
THIS SYSTEM. ARE YOU SURE YOU WANT TO PROCEED? (y or n): y
```

## reset reporting\_data

The **reset reporting\_data** command resets all reporting data.

### Syntax

```
reset reporting_data
```

This command has no arguments.

### Example

```
Infoblox > reset reporting_data  
WARNING: THIS WILL RESET ALL REPORTING DATA. DO YOU WANT TO PROCEED? (y or  
n): y  
ARE YOU SURE YOU WANT TO PROCEED? (y or n): y
```

## reset snmp

This command clears obsolete SNMP credentials (community strings) of devices polled by Network Insight. After that, Network Insight reguesses the SNMP credentials for each device.

You can run this command on the following Grid members:

- Grid Master
- Discovery Consolidator
- Discovery Probes

When the command is executed on the Grid Master or discovery Consolidator, the SNMP credentials are reset for all devices on all discovery Probes. When the command is executed on a discovery Probe, the SNMP credentials are reset for devices only on this Probe.

## Syntax

```
reset snmp
```

This command has no arguments.

## Example

```
Infoblox > reset snmp
```

```
Resetting SNMP credentials will clear all assigned SNMP credentials and cause  
to reguess them for each device.
```

```
Do you wish to continue to the SNMP credentials (y or n): y
```

```
Probe probe.blox: Success
```

```
The command 'reset snmp' completed successfully.
```

```
Infoblox >
```

## reset ssh\_keys

The **reset ssh\_keys** command resets the SSH keys of the system.

## Syntax

```
reset ssh_keys
```

This command has no arguments.

## Example

```
Infoblox > reset ssh_keys
```

```
The system's SSH keys were reset.
```

## restart service

Use the **restart service** command to restart services on a member. You can start individual service on the member, provided that the service is enabled. Note that you can use this command to restart services only on single independent appliances and the active nodes of HA pairs. You cannot use this command on the Grid Master.

## Syntax

```
restart service [dhcp | dns | tftp | http-fd | ftp | ntp | bloxTools | captive_portal]
```

Argument	Description
dhcp	Restart the DHCP service
dns	Restart the DNS service
tftp	Restart the TFTP service
http-fd	Restart the HTTP file distribution service
ftp	Restart the FTP service
ntp	Restart the NTP service
bloxTools	Restart the bloxTools service
captive_portal	Restart the captive portal service

## Examples

```
Infoblox > restart service dhcp
```

## rotate log

The **rotate log** command rolls, or rotates, specified log files. When the audit log, syslog file, and IF-MAP log each reaches its maximum size, the NIOS appliance automatically writes the file into a new file by adding a .0 extension to the first file and incrementing subsequent file extensions by 1. The maximum file size is 100 MB for the audit log, 300 MB for the syslog file, and 120 MB for the IF-MAP log.

Files are compressed during the rotation process, adding a .gz extension following the numerical increment (*file.#.gz*). The first file starts with .0 and subsequent file extensions are incremented by one until it reaches nine. For example, the current log file moves to *file.0.gz*, the previous *file.0.gz* moves to *file.1.gz*, and so on through *file.9.gz*. A maximum of 10 log files (0-9) are kept. When the eleventh file is started, the last log file (*file.9.gz*) is deleted, and subsequent files are renumbered accordingly.

When the debug log file reaches its maximum size, which is 300 MB, the appliance rotates it, but does not compress it. The appliance retains only one previous debug log file to which it adds a .old extension. This command is useful for diagnostic purposes. To export a file to the management system for viewing, you can include it in the support bundle.

To download the support bundle:

1. From the **Grid** tab or **System** tab, select the **Grid Manager** tab or **System Manager** tab, and then click **Download-> Support Bundle** from the Toolbar.
2. Select all options to include configuration and core file information in the output file, then save the tar file to a secure location on the management system.

## Syntax

```
rotate log {syslog | debug | audit}
```

```
rotate file groupname filename [filename2, filename3, ...]
```

Argument	Description
syslog	Syslog file
debug	Debug log file
audit	Audit log file

## Examples

```
Infoblox > rotate log debug
```

```
The selected log file has been rotated to infoblox.log.0.gz
```

```
Infoblox > rotate log audit
```

```
The selected log file has been rotated to audit.log.0.gz
```

## set admin\_group\_acl

Use the set admin\_group\_acl disable command to disable ACL settings for a specific admin group. You will receive an error message when you try to disable a non-existing admin group.

## Syntax

```
set admin_group_acl disable <Admin Group name>
```

Argument	Description
name	Disables ACL settings for a specific admin group.

## Examples

```
Infoblox > set admin_group_acl disable some group
```

```
ACL setting for 'some group' was disabled.
```

```
Infoblox > set admin_group_acl disable nonexistinggroup
```

```
Invalid name.
```

## set adp

The **set adp** command enables or disables ADP (Advanced DNS Protection) on the supported platform. You can use this command only if **Threat Protection** (hardware based) or **Threat Protection (software add-on)** licenses are installed on the platform.

## Syntax

```
set adp
```

### Commands for Threat Protection (software add-on):

- `set adp log <level>`: Use this command to set the threat protection log level, where log level is between **0 (emergency)** and **6 (info)**. The default value is **6 (info)**.
- `set adp log <emergency|alert|critical|error|warning|notice|info>`: Use this command to set the threat protection log level.
- `set adp monitor-mode <on|off>`: Use this command to enable or disable the threat protection monitor mode on the respective member. The default value is **off**.

## Syntax

```
set adp log <level>
```

```
set adp log <emergency|alert|critical|error|warning|notice|info> set adp  
monitor-mode <on|off>
```

### Commands for Threat Protection (hardware-based):

- `set adp log <level>`: Use this command to set the threat protection log level, where log level is between **0 (emergency)** and **7 (debug)**. The default value is **6 (info)**.
- `set adp log <emergency|alert|critical|error|warning|notice|info|debug>`: Use this command to set the threat protection log level.



- `set adp monitor-mode <on|off>` : Use this command to enable or disable the threat protection monitor mode on the respective member. The default value is **off**.

## Syntax

```
set adp log <level>
set adp log <emergency|alert|critical|error|warning|notice|info|debug> set
adp monitor-mode <on|off>
```

## set allow\_query\_domain

You can use the `set allow_query_domain` command to add, update, or delete an allow query domain ACL for the domain of a DNS view. The allow query domain is an ACL that allows or denies the client's request for query access to a domain. To view the list of all domains that have allow query domain ACLs, see [show allow\\_query\\_domain](#). To view the list of DNS views that have allow query domain ACLs, see [show allow\\_query\\_domain\\_views](#).

### Note

You can set the allow query domain named ACL for the domain of a DNS view on the Grid Manager, and all the Grid members will inherit the setting. However, make sure that you do not set the allow query domain named ACL on parental control domain.

## Limitations

- The allow query domain ACL is a named ACL and the `set allow_query_domain` CLI accepts a single named ACL per domain.
- The allow query domain ACL supports BIND. It does not support Unbound and DNS Cache Acceleration. Hence, DNS Cache Acceleration caching is bypassed for domains that are listed in the allow query domain ACL.

## Syntax

```
set allow_query_domain add [view-name] <domain-name> <named-acl>
set allow_query_domain update [view-name] <domain-name> <named-acl>
set allow_query_domain delete <view-name> [domain-name]
```

Argument	Description
<code>add</code>	Adds an ACL for allow query domain for the domain of a DNS view.
<code>update</code>	Updates an ACL for allow query domain for the domain of a DNS view.
<code>delete</code>	Deletes an ACL from allow query domain for the domain of a DNS view. To delete an ACL from the allow query domain, provide the view name of the domain.

Argument	Description
<code>view-name</code>	View in which ACL is configured.
<code>domain-name</code>	Domain to which the ACL is applied.
<code>named-acl</code>	The named ACL defined.

## Examples

To add an ACL, complete the following:

```
Infoblox > set allow_query_domain add default foo.com Named_ACL
```

```
Adding ACL 'Named_ACL' for domain 'foo.com' under view 'default'
Added successfully.
```

A DNS service restart is required in order for the configured ACLs to take effect

To update an ACL, complete the following:

```
Infoblox > set allow_query_domain update default foo.com Named_ACL2
```

```
Updating ACL 'Named_ACL2' for domain 'foo.com' under view 'default'
Updated successfully.
```

A DNS service restart is required in order for the configured ACLs to take effect

To delete an ACL, complete the following:

```
Infoblox > set allow_query_domain delete default foo.com
```

```
Deleting ACL of domain 'foo.com' from view 'default'
Deleted successfully.
```

A DNS service restart is required in order for the configured ACLs to take effect

## set apache\_https\_cert

Use the `set apache_https_cert` command to select one of the previously uploaded HTTPS certificates. NIOS displays the current certificate and all the previously uploaded certificates. You must choose the certificate that you want to use. The current certificate is then replaced with the certificate that you choose.

## Syntax

```
set apache_https_cert
```

## Example

```
Infoblox > set apache_https_cert
Current apache certificate:
  Serial: 7976560e71f701e1a7ee7865fe87d5a4
  Common name: ib-10-34-128-114.infoblox.com
Available certificates:
  1. Serial: 0c8af1b24b1f58bb3d0d05e159841656 , Common name: www.infoblox.com
  2. Serial: 4a73ac27c92a3f731696c1ec0874143d , Common name:
ib-10-34-128-114.infoblox.com
  3. Serial: 26a52734a316c30d43e30b66a6782b18 , Common name:
ib-10-34-128-114.infoblox.com
  4. Serial: 0720ccf94062234db372dd4c8df39dbb , Common name:
ib-10-34-128-114.infoblox.com
  5. Serial: 6bb99aedde38bfe1e1402aa19507a0e1 , Common name:
ib-10-34-128-114.infoblox.com
  6. Serial: 1dc7624dd221e1900aae0e1eec97fb59 , Common name:
ib-10-34-128-114.infoblox.com
  7. Serial: 7976560e71f701e1a7ee7865fe87d5a4 , Common name:
ib-10-34-128-114.infoblox.com

Select certificate (1-7) or q to quit: 2
Are you sure you want to do this? (y or n): y
Certificate updated
```

## set auto\_provision

The **set auto\_provision** command enables and disables auto-provisioning for the NIOS appliance. You cannot enable auto-provisioning for an appliance if a static IP address is already set for an appliance. Note that auto-provisioning can be enabled only on single appliances. To view the status of auto-provisioning for a NIOS appliance, see [show auto\\_provision](#).

## Syntax

```
set auto_provision {on | off}
```

Argument	Description
on	Enables auto-provisioning on an appliance.
off	Disables auto-provisioning on an appliance.

## Examples

Turn on auto-provisioning on an appliance

```
Infoblox > set auto_provision on
```

Turn off auto-provisioning on an appliance

```
Infoblox > set auto_provision off
```

## set bfd

You can use the `set bfd` command to set the BFD logging level. The default logging level is 'informational'. Changing the BFD logging level might cause disruption in advertising due to `bfdd.conf` change. To view the BFD session details, see [show bfd](#).

## Syntax

```
set bfd log [ debugging | informational | notifications | warnings | errors  
| critical | alerts | emergencies ]
```

This command has no arguments.

## Example

```
Infoblox > set bfd log debugging
```

## set bgp log

The `set bgp log` command sets the verbosity level of the BGP routing services and writes statistical information to the syslog. The information in syslog can be helpful for diagnostic purposes. When viewing the syslog file, lines with names such as `bgp statistics` are the BGP statistical information. To view information about the BGP protocol running on the member, see [show bgp](#).

### Note

To use this command, the NSQ software package must be installed.

## Syntax

```
set bgp log {debugging | informational | notifications | warnings | errors |  
critical | alerts | emergencies }
```

Argument	Description
debugging	The verbosity level at which you select to write BGP statistics to syslog.
informational	
notifications	
warnings	
errors	
critical	
alerts	
emergencies	

## Example

```
Infoblox > set bgp log debugging
```

## set bloxtools

You can use the `set bloxtools` command to permanently remove the bloxTools environment and all its data from the appliance. You can also use this command to clear only the user uploaded data and reset the bloxTools environment to the factory default.

### WARNING

When you use this command, bloxTools data is permanently removed from the appliance.

To view the bloxTools status, see [show bloxtools](#).

You can download a copy of the existing bloxtools data using an FTP or SFTP client before you use this command to permanently delete the data. For information, refer to the *Infoblox NIOS Administrator Guide*.

### Note

bloxTools data files are not automatically removed when the bloxTools environment is disabled. You must use this command to manually delete bloxTools data.

## Syntax

```
set bloxtools reset {all | data}
```

Argument	Description
all	Clears all bloxtools related files, which include both the system and data files. You can use this argument to remove the entire bloxtools environment and its data from your appliance.
data	Clears only the user uploaded data that is related to bloxTools and reset the bloxTools environment to the factory default.

## Examples

### Delete all bloxtools data files

```
Infoblox > set bloxtools reset all
```

```
This will erase all Bloxtools data permanently. Do you want to proceed? (y or n):y
```

```
Are you sure you want to do this (y or n): Bloxtools reset.
```

### Delete bloxtools user data

```
Infoblox > set bloxtools reset data
```

```
This will erase all Bloxtools data permanently. Do you want to proceed? (y or n):y
```

```
Are you sure you want to do this (y or n): Bloxtools reset.
```

## set cc\_mode

You can use the `set cc_mode` command to set the Common Criteria mode. To enable or disable Common Criteria configuration, connect to the CLI console, and then enter the `set cc_mode` command. This command will restart the system when it exits the Common Criteria mode. If the system is enabled for Common Criteria, it will reboot in order to go through boot time self tests. You can use this command only on the Grid Master. The setting is propagated to all Grid members. You must restart the members after the configuration change. You can use the `reset all` command to clear the Common Criteria mode. For information about `reset all`, see [reset all](#).



### Note

Factory reset must be done before using the Common Criteria mode.

## Syntax

```
set cc_mode
```

This command has no arguments.

## Examples

```
Infoblox > set cc_mode
Enable Common Criteria mode (grid-level)? (y or n): y New Common Criteria
Mode Settings:
Common Criteria mode enabled: Yes is this correct? (y or n): y
Please refer to the Guidance Documentation Supplement Appendix of the
NIOS Administrator Guide for the requirements to operate a grid in a common
criteria compliant manner.
The system will be rebooted to place it into common criteria mode. Are you
sure you want to continue (y or n): y
[ ]
SYSTEM REBOOTING!

Connection to 10.35.111.3 closed.
```

## set certificate\_auth\_admins

Use the `set certificate_auth_admins` command to disable the certificate authentication service and allow users to log in without validation. Note that when you disable the certificate authentication service, the appliance terminates administrative sessions for all admin users.

## Syntax

```
set certificate_auth_admins disable username
```

Argument	Description
username	Disables certificate authentication service and allows users to log in without validation.

## Examples

```
Infoblox > set certificate_auth_admins disable admin
Certificate authentication for 'admin' was disabled.
```

## set certificate\_auth\_services

Use the `set certificate_auth_services disable name` command to disable a specific certificate authentication service. You will receive an error message when you try to disable a non-existing certificate authentication service.

### Syntax

```
set certificate_auth_services disable name
```

Argument	Description
name	Disables specified certificate authentication service.

### Examples

```
Infoblox > set certificate_auth_services disable name
```

```
Certificate authentication for 'name' was disabled.
```

```
Infoblox > set certificate_auth_services disable DoD CaC
```

```
Certificate Authentication Service for 'DoD CaC' was disabled.
```

```
Infoblox > set certificate_auth_services disable Some Name
```

```
Invalid Name.
```

## set check\_auth\_ns

The `set check_auth_ns` command enables/disables new functionality of checking NS RRset in a response's authority section before overriding delegation NS RRset in recursive cache.

### Syntax

```
set check_auth_ns <true|false>
```

Argument	Description
true	Enables new functionality of checking NS RRset in a response's authority section.
false	Disables new functionality of checking NS RRset in a response's authority section.

### Examples

```
Infoblox > set check_auth_ns true
```



## set circ\_txn\_trace

The `set circ_txn_trace` command enables or disables a circular trace buffer for DB transaction. The file used for capturing the trace buffer is of fixed size, so tracing can be left running for an indefinite period without considering the disk space when trying to capture a problem.

### Syntax

```
set circ_txn_trace [on|off]
```

Argument	Description
on	Enable tracing for circular DB transactions
off	Disable tracing for circular DB transactions

### Example

```
Infoblox > set circ_txn_trace on
```

```
Infoblox > set circ_txn_trace off
```

```
Circular txn trace generated /storage/cores/circ.1632251896.gz
```

## set cloud\_services\_portal\_force\_refresh

The `set cloud_services_portal_force_refresh` command set the flag to request all domains detected by Infoblox Threat Analytics engine in the Cloud.

### Syntax

```
set cloud_services_portal_force_refresh
```

There are no arguments for this command.

### Examples

```
Infoblox > set cloud_services_portal_force_refresh  
Do you want to proceed? (y or n):n
```

## set connection\_limit

You can use the set **connection\_limit** command to set the per client IP address maximum connection limit for the following protocols: HTTP and HTTPS. Note that maximum connections here refer to the network level connections, not application level connections. For example, an HTTPS connection limit of 4 means that there can be a maximum of four TCP connections between any given client IP address and the appliance using the HTTPS protocol. Valid values are from 0 to 2147483647, where 0 means no limit. The default value is 20 for all protocols.

### Note

Setting a low connection limit may have a negative effect on client functionality. For example, some versions of the Firefox browser require at least four TCP connections to function correctly with the appliance. Setting an HTTPS connection limit below four may result in certain browser issues.

To view the current connection limit, see [show connection\\_limit](#).

## Syntax

```
set connection_limit {http | https}
```

Argument	Description
http	Setting maximum connection limit for the HTTP protocol. Valid values are from 0 to 2147483647. The default value is 20.
https	Setting maximum connection limit for the HTTPS protocol. Valid values are from 0 to 2147483647. The default value is 20.

## Examples

### Set the Per Client Address Maximum Connection Limit for the HTTP Protocol

```
Infoblox > set connection_limit http 150
```

## set debug

The **set debug** command enables and disables debug logging for the NIOS appliance. Debug logging is the most extensive and verbose logging that is available on the appliance. It captures all levels of messaging. The output is written into the debug log file. For information on how to view this output, see [show debug](#).

Use this command to capture specific occurrences. However, only use it for short periods of time. Do not leave it running for extended periods of time. Due to the amount of data that is captured, leaving this feature running for any length of time can affect the performance of the appliance. For this reason, it is best to use this command during non-peak hours.

### Note

Infoblox recommends that you turn debug logging off, unless Infoblox Support specifically directs you to turn this feature on. If you leave debug logging turned on, it can cause performance issues.

## Syntax

```
set debug [distribution|upgrade|firewall|ntp|slog|all] [on|off]
```

Argument	Description
all	Specifies debug logging for all services as enabled or disabled.
on	Enables debug logging.
off	Disables debug logging.

## Examples

### Enable debugging

```
Infoblox > set debug all on  
Enabled debug logging for : all
```

### Disable debugging

```
Infoblox > set debug all off  
Disabled debug logging for: all
```

## set debug\_analytics

The `set debug_analytics` command enables or disables debugging of Analytics service.

## Syntax

```
set debug_analytics [on|off]
```

Argument	Description
on	Enables debugging of Analytics service
off	Disables debugging of Analytics service

## Examples

### Enable debugging

```
Infoblox > set debug_analytics on
```

### Disable debugging

```
Infoblox > set debug_analytics off
```

## set default\_revert\_window

Use the `set default_revert_window` command to configure the Grid default time window for reverting a member after it was upgraded from NIOS 6.4.0 to a later release. Note that you can only change the default value on the Grid Master. When you change the default value, the new revert window affects only the members that have not been upgraded.

## Syntax

```
set default_revert_window hours
```

Argument	Description
<i>hours</i>	The number of hours configured for the default revert window. The minimum value is 1 and the maximum is 48. The default is 24.

## Example

```
Infoblox > set default_revert_window 36
```

```
Member revert window is currently: 24h
```

```
Member Revert Window being changed to 36 hours Is this correct? (y or n): y
```

```
Member Revert Window change will only affect members which are not yet  
upgraded. Member Revert Window is changed.
```

## set default\_route

The `set default_route` command allows you to configure the default gateway for the NIOS appliance. You can set the gateway address of LAN1 or LAN2 as the default route. You can also specify the IPv4 or IPv6 gateway address. You can also set an optional VLAN gateway address and make it the default route.

## Syntax

```
set default_route LAN1|LAN2
```

```
set default_route IPv4gateway [IPv6gateway] | IPv6gateway [IPv4gateway]
```

Argument	Description
LAN1 LAN2	Specifies the LAN1 gateway address. Specifies the LAN2 gateway address.
IPv4gateway IPv6gateway	Specifies the IPv4 gateway address. Specifies the IPv6 gateway address.

## Example

```
Infoblox > set default_route LAN1
```

```
Infoblox >set default_route LAN2
```

```
Infoblox > set default_route 10.35.0.1 2620:10a:6000:2400::1
```

```
Infoblox > set default_route 2620:10a:6000:2400::1
```

## set delete\_tasks\_interval

Use the **set delete\_tasks\_interval** command to configure the time interval the appliance waits until it deletes completed and rejected tasks from the system. Grid Manager displays these tasks in the **Task Manager** tab until they are deleted from the system. By default, Grid Manager displays these tasks for 14 days. You can configure this time interval. Valid values are from 1 to 30 days.

Use the **show delete\_tasks\_interval** command to view the current time interval. For information, see [show delete\\_tasks\\_interval](#).

## Syntax

```
set delete_tasks_interval days
```

Argument	Description
days	The number of days completed and rejected tasks are displayed in the <b>TaskManager</b> tab before they are deleted. The minimum value is 1 and the maximum is 30. The default is 14.

## Example

```
Infoblox > set delete_tasks_interval 25
```

```
Current delete tasks interval is 14 days
```

```
The delete tasks interval has been changed to 25 days Is this correct? (y or n): y
```

```
The delete tasks interval has been changed.
```

## set dhcpd\_recv\_sock\_buf\_size

You can use the **set dhcpd\_recv\_sock\_buf\_size** command to tune the DHCP receive socket buffer memory. The DHCP receive socket buffer holds DHCP packets that are queued on the UDP (User Datagram Protocol) port from the NIC (Network Interface Controller). This command is useful when you want to increase the DHCP receive buffer size to accommodate occasional burst traffic and high volume DHCP requests. Use the [show dhcpd\\_recv\\_sock\\_buf\\_size](#) to view the current buffer size.

### Note

Ensure that you use this command only when you are dealing with burst traffic situations in high volume deployments.

## Syntax

```
set dhcpd_recv_sock_buf_size N [120 <= N <= 8192, 1536=default]
```

## Examples

Argument	Description
<i>N</i>	The number of kilobytes to which you want to set the BIND receive socket buffer size. The minimum is 120 kilobytes and the maximum is 8192. The default is 1536.

```
Infoblox > set dhcpd_recv_sock_buf_size 1500
```

```
DHCP service restart is required in order for the changed value to take effect
```

## set extra\_dns\_name\_validations

Use the **set extra\_dns\_name\_validations** command to enable or disable additional validation on host names that takes place at the time of data entry when creating zones, subzones, and records of type A, AAAA, host record, ALIAS, CAA, MX, and NS. The additional validation invokes the same functions in the BIND library that the DNS service invokes when it is loading a zone and this can impact the performance of your current operation.

By default, the additional validation is disabled. When the validation is enabled, if you enter an invalid name, NIOS displays an error message and does not save the record until the name that you enter passes the additional validation.

## Syntax

```
set extra_dns_name_validations [on|off]
```

Argument	Description
on	Enables the additional DNS name validation that takes place when you create a zone or a resource record.
off	Disables the additional DNS name validation that takes place when you create a zone or a resource record.

## Example

```
Infoblox > set extra_dns_name_validations on
WARNING: This has a performance impact when creating or importing multiple
records at a time. Do you want to continue? (y or n): y
Extra DNS name validations is turned on
```

## set disable\_gui\_one\_click\_support

The **set disable\_gui\_one\_click\_support** command permanently disables the feature to submit technical support requests through the Infoblox GUI.

## Syntax

```
set disable_gui_one_click_support
```

There are no arguments for this command.

## Example

```
Infoblox > set disable_gui_one_click_support
WARNING: Once you permanently disable this feature, you cannot enable it
again.
Are you sure you want to proceed? (y or n): y
```

## set disable\_https\_cert\_regeneration

Use the **set disable\_https\_cert\_regeneration** command to turn on or off the automatic regeneration of a self-signed HTTPS certificate.

NIOS regenerates a certificate in the following scenarios:

- If you change a host name and the new name does not match the name of the existing certificate. This is especially useful for wildcard certificates.

- If the certificate is self-signed and the regeneration is enabled (it is enabled by default) , restarting NIOS or changing the host name or IP address causes NIOS to regenerate the certificate.

## Syntax

```
set disable_https_cert_regeneration [on|off]
```

Argument	Description
on	Disables the automatic regeneration of the self-signed HTTPS certificate. Automatic certificate regeneration is enabled by default.
off	Enables the automatic regeneration of the self-signed HTTPS certificate

## Example

```
Infoblox > set disable_https_cert_regeneration off
HTTPS certificate regeneration enabled.
```

## set dns

The **set dns** command enables you to control the DNS cache. You can flush the cache file of a DNS view or flush a particular entry from a cache file. You can also flush a specific domain and its subdomains from the DNS cache. In addition, you can schedule an inbound zone transfer from an external primary server. This command displays IDN data in punycode. You can also use this command to delete cache files from the default DNS view for DNS cache acceleration on IB-FLEX.

## Syntax

```
set dns flush all [dns_view]
set dns flush name name [dns_view]
set dns transfer zone [dns_view]
set dns flush tree <part-of-domain-name> [dns_view]
```

Argument	Description
all	Flushes the cache file from the default view.
dns_view	Specifies a particular DNS view.
name	Flushes the specific entry from the cache.
zone	Specifies the zone for the inbound transfer from an external primary server.



Argument	Description
<code>tree &lt;part-of-domain-name&gt;</code>	Flushes the specified domain and its subdomains from the DNS cache. For example, if you enter the domain name <code>corp100.com</code> , then the specified domain and its subdomains such as <code>www.corp100.com</code> , <code>corp100.com</code> , <code>x.corp100.com</code> , etc. are cleared from the DNS cache.

## Example

Flush the cache file from the default DNS view

```
Infoblox > set dns flush all
```

Flush the specified domain and its subdomains from the default DNS view

```
Infoblox > set dns flush tree corp100.com default
```

## set dns-accel

The `set dns-accel` command enables you to set certain parameters for the **DNS Cache Acceleration** feature. This command is available for:

- IB-FLEX only if the **Flex Grid Activation** license is present in the Grid.

## Syntax

```
set dns-accel log <level>
```

```
set dns-accel log <emergency|alert|critical|error|warning|notice|info|debug>
```

Argument	Description
<code>level</code>	The DNS Cache Acceleration log level, where log level is between 0 and 7, and the default value is 6.
<code>&lt;emergency alert critical error warning notice info debug&gt;</code>	Specifies one of these log levels.

## Example

```
Infoblox > set dns-accel log 2
```

```
Infoblox > set dns-accel log notice
```

## set dns\_rrl

Use the **set dns\_rrl** command to configure RRL (Response Rate Limiting) settings for the Grid or members. Changes made using this command are applied immediately to an active DNS resolver, although there could be replication delays for Grid Master configuration of other members. Use the [show dns\\_rrl](#) to view the current RRL settings.

### Syntax

```
set dns_rrl enable
set dns_rrl disable
set dns_rrl [member <hostname> | view <viewname>] [override|inherit]
set dns_rrl [member <hostname> | view <viewname>] [enable|disable]
set dns_rrl [member <hostname> | view <viewname> | grid]
    [responses_per_second <number>|disable] [window <number>|default]
    [slip <number>|default|disable] [logging enable|disable|default]
    [log_only true|false|default]
```

#### Note

The set dns\_rrl command accepts the member, view, and grid options only on the Grid Master.

Argument	Description
enable	Enable RRL with previously configured values.
disable	Disable RRL
member <hostname>	The FQDN of the Grid member.
view <viewname>	The name of the DNS view.
override inherit	For a Grid member, specify whether you want to override the Grid RRL settings or inherit them from the Grid.
responses_per_second <number>	The number of DNS responses per second for the RRL. Valid values are from 1 to 1000. Although the BIND default is 0, which means there is no limit or RRL is disabled, you cannot set this to 0 in NIOS. Use the "disable" argument to disable this feature. The default is 100.
window <number> default	A rolling window of time (in seconds) within which DNS responses are tracked. Valid values are from 1 to 3600 seconds. The default is 15 seconds.

Argument	Description
<code>slip &lt;number&gt;/default/disable</code>	The number of UDP requests that the appliance skips before answering with a truncated response. For example, if you set the slip number to 2, the appliance responds to every other UDP request. Valid numbers are from 0 to 10. The appliance does not “slip” if the number is set to 0. The default is 2.
<code>logging enable disable default</code>	Enable or disable the logging of RRL events to the “rate-limit” logging category in syslog. This is enabled by default; however, RRL events are logged to the “rate-limit” category only if RRL is enabled. Note that inheritance for logging categories applies; therefore, you must explicitly override Grid logging categories on a member for changes to the member logging setting to take effect. You cannot use the <code>dns_rrl</code> override command to override logging configuration for a member. Changes made to this setting require a service restart.
<code>log_only true false default</code>	Set this to true to test the RRL settings without dropping any requests. Set this to false to enable RRL. The default is false.

## Examples

### Configure the Grid RRL settings on the Grid Master

```
Infoblox > set dns_rrl responses_per_second 100 log_only false window default slip 2 logging disable
```

#### Note

You can also execute the above command on a Grid member to change the RRL settings for that member.

### Override the Grid RRL settings on a Grid member

```
Infoblox > set dns_rrl corp100.com override responses_per_second 300  
log_only false window 200 slip 3
```

### Inherit the Grid RRL settings on a Grid member

```
Infoblox > set dns_rrl corp100.com inherit
```

## set docker\_bridge

The `set docker_bridge` command resets the IP address of the current Docker bridge to the IP address that you specify. You can view the IP address of the current Docker bridge by running the `show docker_bridge` command.

## Syntax

```
set docker_bridge
```

This command has no arguments.

## Example

```
Infoblox > set docker_bridge
Enter Bridge Gateway/CIDR: 172.17.0.1/16
New Docker Bridge settings:
Bridge Gateway/CIDR: 172.17.0.1/16
Current Docker Bridge settings:
Bridge Gateway/CIDR: 172.16.0.1/16
WARNING: This operation will restart the system to reset the current docker
bridge settings.
Do you want to proceed? (y or n):y
Docker bridge settings have been updated.
System will be restarted.
```

### Note

- The warning message to restart the system is displayed only if the **Enable Recursive Queries Forwarding to BloxOne Threat Defense Cloud** checkbox is selected.
- Infoblox recommends that you run the `set docker_bridge` command only on non-routable IP addresses.
- Infoblox does not recommend that you run the `set docker_bridge` command on a passive HA node

## set dscp

Use the `set dscp` command to configure the DSCP value, which determines the PHBs (per-hop behaviors) on DiffServ compliant nodes and enables priorities of services to be assigned to network traffic. When you set the DSCP value, the appliance implements QoS (quality of service) rules based on your configuration so you can effectively classify and manage your critical network traffic. Note that on an appliance, all outgoing IP traffic on all interfaces uses the same DSCP value. You can configure this value for the Grid. You can also override the Grid setting for individual members. DSCP is supported on both IPv4 and IPv6 transports. This feature is currently supported on the following Infoblox appliances: Trinzic 2210, Trinzic 2220, and Infoblox-4010. For information about these appliances, refer to the respective installation guides.

## Syntax

```
set dscp grid [value]
set dscp member [value]
set dscp member inherit
```

Argument	Description
<code>value</code>	The DSCP value. You can enter a value from 0 to 63. The default is 0 and it represents the lowest priority.

## Example

### Set the Grid DSCP value

```
Infoblox > set dscp grid 32
```

### Override the Grid DSCP value for a specific member

```
Infoblox > set dscp member 20
```

### Inherit the Grid DSCP value

```
Infoblox > set dscp member inherit
```

## set enable\_dnstap

Use the `set enable_dnstap` command to enable or disable using dnstap to log DNS queries and responses. For information about using dnstap, see [Capturing DNS Queries and Responses](#). To view the status of the dnstap configuration, see [show dnstap-status](#).

## Syntax

```
set enable dns_tap [on|off]
```

Argument	Description
<code>on</code>	Enables the use of the dnstap log format to log DNS queries and responses at high rates to well-known destinations.
<code>off</code>	Disables the use of the dnstap log format to log DNS queries and responses.

## Examples

To enable the use of the dnstap log format:

```
Infoblox > set enable_dnstap on
```

```
Override Grid settings and configure DNSTAP on member level? (y or n):
```

The input `y` overrides Grid settings and enables member-level dnstap configuration. The input `n` inherits Grid settings; no further configuration is performed.

To disable the use of the dnstap log format:

```
Infoblox > set enable_dnstap off
```

## set enable\_match\_recursive\_only

Use the `set enable_match_recursive_only` command to enable or disable the match-recursive-only option for a specific DNS view on a specific Grid member. You can also use the match-recursive-only option in combination with the Match Clients and Match Destinations settings to restrict and filter client access for specific DNS views on specific Grid members. For information about how to use these features, refer to the *Infoblox NIOS Administrator Guide*. If you want to enable this setting for a DNS view, ensure that the **Enable Recursion** setting is enabled for the DNS view on the specified member.

To check the status of this setting for all DNS views on a Grid member, use the `show enable-match-recursive-only` CLI command. For information, see [show enable\\_match\\_recursive\\_only](#).

## Syntax

```
set enable_match_recursive_only <true|false|inherit> [dns_view]
```

Argument	Description
<code>&lt;true false inherit&gt;</code>	Set the <code>enable_match_recursive_only</code> setting to <code>true</code> , <code>false</code> , or <code>inherit</code> . The default value is <code>inherit</code> . The <code>true</code> setting enables the match-recursive-only option for the specified DNS view on the specific member; <code>false</code> disables it for the specified DNS view on the specific member. <code>Inherit</code> represents the setting for the DNS view ( <code>true</code> or <code>false</code> ) that is populated across all members serving that DNS view. By specifying <code>true</code> or <code>false</code> , you override the <code>inherit</code> setting for the specific member. Specifying <code>inherit</code> restores the inherited setting for the specific member.
<code>dns_view</code>	Optional parameter to specify the DNS view. If this parameter is omitted, the setting affects only the default DNS view. If the specific Grid member does not serve the default DNS view, you will receive an error message by omitting this parameter.

## Example

```
Infoblox > set enable_match_recursive_only true (affects default DNS view only if default DNS view is served by member)
```

```
Infoblox > set enable_match_recursive_only true external
```

```
Infoblox > set enable_match_recursive_only false corp100sales
```

```
Infoblox > set enable_match_recursive_only inherit external
```

## set fips\_mode

You can use the `set fips_mode` command to enable the FIPS mode. This command restarts the appliance to go through the boot time self tests when it exits the FIPS mode.

- In a Grid, you can set the FIPS mode only on the Grid Master. The setting is propagated to all Grid members during the joining process. After the configuration is changed, the members will be restarted.
- You can set the FIPS mode on standalone systems.
- In an HA setup, you can set the FIPS mode only on the standalone Grid Master, and then configure it as a node in the HA pair. Perform the same step for the second node of the HA pair. You cannot change the FIPS mode setting on the HA Grid Master or the HA member.

To enable or disable the FIPS configuration, connect to the CLI console, and then enter the `set fips_mode` command. For more information, see [Enabling / Disabling the FIPS Mode](#).

To clear the FIPS mode, you can use the `reset all` command. For more information, see [reset all](#).

### Note

You must perform a factory reset to reset the appliance to its original factory settings before using the FIPS mode.

## Syntax

```
set fips_mode
```

This command has no arguments.

## Examples

```
Infoblox > set fips_mode
Enable FIPS mode? (y or n): y
New FIPS Mode Settings:
  FIPS mode enabled: Yes
  is this correct? (y or n): y
Please refer to the Guidance Documentation Supplement Appendix of the NIOS
Administrator Guide for the requirements to operate a grid in a FIPS
compliant manner.
The system will be rebooted to place it into FIPS mode. Are you sure you
want to continue (y or n): y
Integrity private key and certificate were generated successfully.
Sign executable files by sha256sum...
```

## set hardware-type

Use the `set hardware-type` command for IB-FLEX only. This command enables you to set a supported virtual appliance as an IB-FLEX. For more information, see [About IB-FLEX](#).

To see if an appliance has been configured as an IB-FLEX, you can use the `show hardware-type` command.

## Syntax

```
set hardware-type
```

This command has no arguments.

## Example

When setting a supported virtual appliance as an IB-FLEX:

```
Infoblox > set hardware-type IB-FLEX
```

## set ibtrap

The `set ibtrap` command allows you to specify whether the appliance sends SNMP notifications (traps) and email notifications to the configured trap receivers and email recipients for the specified event category.

## Syntax

```
set ibtrap [category] snmp [true|false] email [true|false]
```

Argument	Description
category	The event category that triggers the trap and/or email notification. Valid values are: Fan, Bloxtools, Disk, Memory, CPU, MGM, HSM, Login, PowerSupply, FTP, TFTP, HTTP, NTP, DNS, DHCP, RootFS, Database, RAID, HA, MSServer, Backup, Clear, SNMP, LCD, SSH, SerialConsole, ENAT, Network, Cluster, Control, OSPF, OSPF6, IFMAP, BGP, CaptivePortal, DuplicateIP, License, System, Syslog, DiscoveryConflict, Reporting, FDUsage, OCSPResponers, DisconnectedGrid, LBDevice, LDAPServers, RIRSWIP
snmp true   false	Specify <code>true</code> to send SNMP traps. Otherwise, specify <code>false</code> .
email true   false	Specify <code>true</code> to send email notifications. Otherwise, specify <code>false</code> .

## Example

```
Infoblox > set ibtrap FTP snmp true email true
```



## set interface

The

### `set interface`

command allows you to configure the speed and duplex settings of the network interfaces (MGMT, HA, LAN1, and LAN2) on single independent appliances only. You cannot configure the network interface settings of appliances after they join a Grid or become HA pairs.

You can use

### `set`

`interface mgmt` to enable the MGMT port and configure its IP address, netmask, and gateway address. You can configure either IPv4 address, IPv6 address or both for the MGMT interface of the appliance. Once the MGMT port is enabled, you can use the command to configure the speed and duplex settings of the MGMT port. You can also use

### `set`

`interface mgmt off` to disable the MGMT port.

Use the [show interface](#) command to view the interface settings.

## Syntax

```
set interface [lan1|ha|lan2] speed [auto|10M|100M|1000M] duplex {auto|half|full}
set interface mgmt [speed [auto|10M|100M|1000M] duplex {auto|half|full}]
set interface mgmt off
```

Argument	Description
<code>lan1</code> <code>lan2</code> <code>ha</code>	Specifies the LAN1 interface. Specifies the LAN2 interface on the Infoblox-250-A, -550-A, -1050-A, -1550-A, -1552-A, and -2000-A appliances. Specifies the HA interface.
<code>speed</code> <code>auto</code> <code>10M   100M   1000M</code>	Specifies the speed of the incoming line rate in Mbps, or allows the appliance to automatically match the line speed.

Argument	Description
<pre> duplex auto half full </pre>	<p>Specifies the duplex speed:</p> <ul style="list-style-type: none"> <li>Automatically adjusts the speed</li> <li>Sets it at half speed</li> <li>Set it at full speed</li> </ul>
<pre> mgmt </pre>	<p>Specifies the management interface.</p>
<pre> mgmt off </pre>	<p>Disables the management system interface.  <b>Note:</b> If this port is not being used, it should be set to off for security reasons.</p>

## Examples

### Enable and configure IPv4 address for the MGMT interface

```

Infoblox >
set interface mgmt

Enable Management port? (y or n):
y

Enter Management IP address:
10.36.1.157

Enter Management netmask [Default: 255.255.255.0]:
255.255.0.0

Enter Management gateway address [Default: 10.36.0.1]:
Configure Management IPv6 network settings? (y or n): n
Restrict Support and remote console access to MGMT port? (y or n):
n

```

```
Management Port Setting:
Management Port Enabled: true
Management IP address: 10.36.1.157
Management netmask: 255.255.0.0

Management Gateway address: 10.36.0.1
Restrict Support and remote console access to MGMT port:false

Is this correct? [ y or n]:
y

Are you sure? (y or n):
y

The management port settings have been updated.
```

## Enable and configure IPv6 address for the MGMT interface

```
Infoblox > set interface mgmt
Enable Management port? (y or n): y
Enter Management IP address: 2620:010A:6000:2400::6508
Enter Management IPv6 Prefix Length [Default: none]: 64
Enter Management IPv6 gateway address [Default: none]:
2620:010A:6000:2400::0001
Configure Management IPv4 network settings? (y or n): n
Restrict Support and remote console access to MGMT port? (y or n): n

Management IPv6 address: 2620:10a:6000:2400::6508/64
```

```
Management IPv6 Gateway address: 2620:10a:6000:2400::1
Restrict Support and remote console access to MGMT port: false
```

```
Is this correct? (y or n): y
Are you sure? (y or n): y
```

```
The management port settings have been updated
```

Specify the MGMT interface speed after the port is enabled

```
Infoblox >
set interface mgmt speed 10M duplex full

Setting Management interface speed to: 10M and duplex to: full
Is this correct? [ y or n]:
y

The network interface settings have been updated.
```

Specify the LAN interface speed

```
Infoblox >
set interface lan1 speed 10M duplex full

Setting LAN1 interface speed to: 10M and duplex to: full
Is this correct? [ y or n]:
y
```

```
The netwok interface settings have been updated.
```

## Specify the HA interface speed

```
Infoblox >
```

```
set interface ha speed 100M duplex half
```

```
Setting HA interface speed to: 100M and duplex to: half
```

```
Is this correct? [ y or n]:
```

```
y
```

```
The netwok interface settings have been updated.
```



### Note

This command is not supported on vNIOS appliances

## set ip\_rate\_limit

The **set ip\_rate\_limit** commands enable and disable rate limiting UDP traffic from source port 53, configure rate limiting rules that control the traffic, and remove rate limiting rules. Once you enable rate limiting, the current rate limiting rules take effect.

This command is useful when you want to mitigate cache poisoning on your DNS server by limiting the UDP traffic or blocking connections from source port 53.

## Syntax

```
set ip_rate_limit {on | off}  
set ip_rate_limit remove {source all | all | source ip-address[/mask]}  
set ip_rate_limit add source {all | ip_address [/mask]} limit 0  
set ip_rate_limit add source {all | ip_address [/mask]} limit packets/m [burst  
burst_packets]
```

Argument	Description
On	Enables rate limiting from UDP port 53.
Off	Disables rate limiting from UDP port 53.

Argument	Description
<pre>add source   all   ip_address/mask   limit packets   burst <i>burst_packets</i></pre>	<p>Configures the rate limiting rules.</p> <p>Enter all or 0.0.0.0 if you want to limit all traffic from all sources.</p> <p>Enter the IP address, and optionally the netmask, from which you want to limit the UDP traffic on port 53.</p> <p>Enter the number of packets per minute that you want to receive from the source.</p> <p>Optionally, enter burst and the number of packets for burst traffic. Burst is the maximum number of packets accepted.</p>
<pre>remove   source all   all   source <i>ip-address/mask</i></pre>	<p>Removes rate limiting rules from all sources or an existing host on UDP port 53.</p> <p>Removes the rate limiting rule that limits traffic from all sources on UDP port 53.</p> <p>Removes all of the rate limiting rules from all sources on UDP port 53.</p> <p>Removes the existing rules for an existing host.</p>

## Examples

### Turn on rate limiting

```
Infoblox > set ip_rate_limit on
Enabling rate limiting will discard packets and may degrade performance.
Are you sure? (y or n):
```

### Turn off rate limiting

```
Infoblox > set ip_rate_limit off
```

### Block all traffic from host 10.10.1.1

```
Infoblox > set ip_rate_limit add source 10.10.1.1 limit 0
```

### Limit the traffic to five packets per minute from host 10.10.1.2/24, with an allowance for burst of 10 packets

```
Infoblox > set ip_rate_limit add source 10.10.1.2/24 limit 5/m burst 10
```

### Remove the rate limiting rule from host 10.10.1.1/24

```
Infoblox > set ip_rate_limit remove source 10.10.1.1/24
```

## set ipam\_web\_ui

The **set ipam\_web\_ui** command enables and disables Grid Manager on vNIOS appliances on Cisco. For information about Grid Manager, refer to the *Infoblox Administrator Guide*.

### Syntax

```
set ipam_web_ui
```

This command has no arguments.

### Example

```
Infoblox > set ipam_web_ui
```

## set ipv6\_disable\_on\_dad

The **set ipv6\_disable\_on\_dad** command enables or disables IPv6 on an interface if a duplicate IPv6 address is detected.

### Syntax

```
set ipv6_disable_on_dad {on | off}
```

Argument	Description
on	Enables IPv6 on an interface.
off	Disables IPv6 on an interface.

### Examples

#### Turn on IPv6 on an interface

```
Infoblox > set ipv6_disable_on_dad on
WARNING: This operation will reboot the system.
Do you want to proceed? (y or n): y
SYSTEM REBOOTING!
```

```
Infoblox > set ipv6_disable_on_dad on
Already on, nothing do be done
```

## Turn off Pv6 on an interface

```
Infoblox > set ipv6_disable_on_dad off
WARNING: This operation will reboot the system.
Do you want to proceed? (y or n): y
SYSTEM REBOOTING!
```

## set ipv6\_neighbor

The `set ipv6_neighbor` command enables definition of an IPv6 neighbor for any of the following: LAN1, LAN2 or MGMT. `Set ipv6_neighbor` also allows deletion of an existing IPv6 neighbor entry on the specified interface. For adding a new neighbor entry, the second required argument is for the link-local MAC address ID of the neighboring interface for the specified LAN/LAN2/MGMT port. Another form of this command allows the flushing of specific or general IPv6 neighbor values from the specified interface. Prefixes and polled neighbor states can also be specified and combined in a statement.

## Syntax

```
set ipv6_neighbor {add|clear} {LAN|LAN2|MGMT} [all] [prefix] ipv6-address
ll_address [state]
```

Argument	Description
<i>ipv6_address</i>	The IPv6 address of the neighboring interface.
<i>ll_address</i>	The 48-bit link-local MAC ID of the neighboring interface. Argument is required for addition of a new IPv6 neighbor entry for the interface.
[all]	Optional argument to clear the entire list of IPv6 entries for the specified interface.
[prefix]	Optional argument needed if all entries are to be flushed for an IPv6 prefix. CIDR mask is required as part of the address specification.
[state]	Optional argument needed if entries of a specific type are to be flushed or defined for an IPv6 prefix. Permitted values for the <b>state</b> argument include the following: <b>permanent</b> , <b>noarp</b> , <b>reachable</b> , and <b>stale</b> .

## Example

```
Infoblox > set ipv6_neighbor add LAN 2001:db8::256:180:c223:214e
02:80:C2:03:DE:05
Infoblox > set ipv6_neighbor add LAN 2001:db8::256:180:c223:214e
```



```
02:80:C2:03:DE:05 permanent
```

```
Infoblox > set ipv6_neighbor clear LAN2 2001:db8::256:180:c223:214e
```

```
02:80:C2:03:DE:05
```

```
Infoblox > set ipv6_neighbor clear LAN prefix 2001:db8:12:256::/64 stale
```

```
Infoblox > set ipv6_neighbor clear LAN all
```

## set ipv6\_ospf

The `set ipv6_ospf` command writes statistical information to syslog. This command provides informational data that can be helpful for diagnostic purposes. Setting the log level for OSPFv3 is the only configuration that can be done for the routing protocol in the NIOS CLI. The statistical information is written (dumped) to syslog. When viewing the syslog file, lines with names such as `ipv6_ospf statistics` are the OSPF statistical information. Use the `show ipv6_ospf` command to view the OSPF settings.

Syslog level describes the types of messages that are sent to syslog. You can identify the syslog information by using the `level` option.

### Note

To use this command, the NSQ software package must be installed.

## Syntax

```
set ipv6_ospf log {level}
```

Argument	Description
<code>level</code>	Writes OSPF statistics to syslog with a specific associated level. The supported log levels are: debugging, informational, notifications, warnings, errors, critical, alerts, and emergencies.

## Example

```
Infoblox > set ipv6_ospf log alerts
```

## set ipv6\_status

The `set ipv6_status` command enables or disables IPv6 on all interfaces. This is not a permanent enable or disable. If you restart your system, IPv6 is enabled once again.

## Syntax

```
set ipv6_status {enable / disable}
```

Argument	Description
enable	Enables IPv6 on all interfaces.
disable	Disables IPv6 on all interfaces.

## Example

```
Infoblox > set ipv6_status enable
```

```
WARNING: This operation will restart the product
```

```
Do you want to proceed? (y or n):y
```

```
Infoblox > show interface
```

```
MGMT:
```

```
IP Address: 10.36.111.3      MAC Address: 00:0C:29:70:D5:F5
```

```
Mask: 255.255.0.0          Broadcast: 10.36.255.255
```

```
MTU: 1500                  Metric: 1
```

```
IPv6 Link: fe80::20c:29ff:fe70:d5f5/64
```

```
IPv6 Status: Enabled
```

```
Negotiation: Disabled
```

```
Speed: 1000M               Duplex: Full
```

```
Status: UP BROADCAST RUNNING MULTICAST
```

```
Statistics Information
```

```
Received
```

```
packets: 25                bytes: 1518 (1.4 KiB)
```

```
errors: 0                  dropped: 0
```

```
overruns: 0                frame: 0
```

```
Transmitted
```

```
packets: 3                 bytes: 218 (218.0 b)
```

```
errors: 0                  dropped: 0
```

```
overruns: 0                carrier: 0
```

```
Collisions: 0 Txqueuelen: 1000
```

```
Infoblox > set ipv6_status disable
```

```
WARNING: This operation will disable IPv6 communication
```

```
Do you want to proceed? (y or n):y
```

```
Infoblox > show interface
```

```
MGMT:
```

```
IP Address: 10.36.111.3 MAC Address: 00:0C:29:70:D5:F5
```

```
Mask: 255.255.0.0 Broadcast: 10.36.255.255
```

```
MTU: 1500 Metric: 1
```

```
IPv6 Link:
```

```
IPv6 Status: Disabled
```

```
Negotiation: Disabled
```

```
Speed: 1000M Duplex: Full
```

```
Status: UP BROADCAST RUNNING MULTICAST
```

```
Statistics Information
```

```
Received
```

```
packets: 606 bytes: 66780 (65.2 KiB)
```

```
errors: 0 dropped: 0
```

```
overruns: 0 frame: 0
```

```
Transmitted
```

```
packets: 10 bytes: 540 (540.0 b)
```

```
errors: 0 dropped: 0
```

```
overruns: 0 carrier: 0
```

```
Collisions: 0 Txqueuelen: 1000
```

## set lcd keys or set lcd

The set lcd keys or set lcd command enables and disables the LCD input keys. Turning off the LCD input keys prevents anyone from manually changing the IP address on the NIOS appliance. Infoblox recommends this practice as a security measure for remote appliances.

## Syntax

```
set lcd keys {off | on}
```

Argument	Description
<code>off</code>	Disables the LCD input keys on the appliance.
<code>on</code>	Re-enables the LCD input keys on the appliance.

## Examples

### Disable the lcd keys

```
Infoblox > set lcd keys
```

### Enable the lcd keys

```
Infoblox > set lcd keys on
```

```
Turning ON the LCD display...
```

#### Note

You cannot enable or disable the LCD input keys on vNIOs appliances. You can configure the LCD input keys only on a Grid Master. On a vNIOs appliance, the `set lcd keys` or `set lcd` command generates an error.

## set lcd\_settings

The `set lcd_settings` command enables you to set the display settings of an LCD. You can specify the number of seconds after which the LCD screen must reduce the brightness if there is no keypad activity and specify the brightness level. You can also use this command to set the UID (unit identification) button on Trinzic appliances. For more information, see [set lcd\\_settings hwident](#).

## Syntax

Argument	Description
<code>&lt;seconds&gt;</code>	Sets the number of seconds after which the LCD screen should automatically dim. The auto-dim value should be in the range of 5 to 3600.
<code>&lt;level&gt;</code>	Sets the brightness of the LCD screen. Brightness levels are from 1 to 10.

## Example

```
Infoblox > set lcd_settings autodim 8
```

```
Infoblox > set lcd_settings brightness 5
```

## set lcd\_settings hwident

The `set lcd_settings hwident` command enables and disables the UID (unit identification) button on Trinzic appliances. When you enable the UID button, the LCD panel on the front panel blinks and the UID LED on the rear panel glows blue. In a rack environment, the UID feature allows you to easily identify the appliance when moving between the front and rear of the rack.

## Syntax

```
set lcd_settings hwident {off | on}
```

Argument	Description
off	Disables the UID feature on the Trinzic appliance.
on	Enables the UID feature on the Trinzic appliance.

## Examples

### Disable the UID feature

```
Infoblox > set lcd_settings hwident off
```

```
Turning OFF the UID feature
```

### Enable the UID feature

```
Infoblox > set lcd_settings hwident on
```

```
Turning ON the UID feature
```



#### Note

You cannot enable or disable the UID feature on vNIOs appliances. You can configure the UID feature only on Trinzic appliances. On a vNIOs appliance, the `set lcd_settings hwident` command generates an error.

## set license

The `set license` command installs a license upon entering a valid license string. You must send an email request to Infoblox to receive a unique license string for your NIOS appliance. Copy the string directly from the email, and then use CTRL + V to insert it after the CLI command prompt. Use the `show license` command to view the license settings. This command is used to install both static (per member) and Grid-wide licenses.

### Note

You can install a temporary 60-day license that allows your system to be fully functional while waiting to receive your permanent license. For more information, see [set temp\\_license](#).

## Syntax

```
set license
```

This command has no arguments.

## Example

```
Infoblox > set license
Enter license string: EQAAAKS4n90WFGNUSirwvyUT9/z
Install license? (y or n): y
Infoblox > set license
Enter license string: HQAAALsak0zDKirMdaUsG2Yfk/j0BkhoFjhVfEtu36dJ
Install license? (y or n): y

License (grid-wide) is installed.

The UI needs to be restarted in order to reflect this license change.

Restart UI now, this will log out all UI users? (y or n):y
```

## set lines

The `set lines` command specifies the number of lines that the appliance displays when you execute a `show` command during a session. The default is 20 lines. You can also configure permanent page settings or enter zero (0) to set paging off.

## Syntax

```
set lines [num | permanent]
```

Argument	Description
num	The number of lines the appliance displays when you execute a <b>show</b> command.
permanent	Configures permanent page settings.

## Examples

Set the number of lines displayed on each page to 4:

```
Infoblox > set lines 4
Number of scroll lines set to 4.
Infoblox > show log
May 31 13:30:05 (none) syslog-ng[892]: syslog-ng version 1.6.11 starting
May 31 13:30:05 (none) kernel: Linux version 2.6.17.4 (root@build-aslan)
(gcc version
3.2.1) #1 SMP Fri May 18 19:44:21 EDT 2013
May 31 13:30:05 (none) kernel: BIOS-provided physical RAM map:
May 31 13:30:05 (none) kernel: BIOS-e820: 0000000000000000 -
00000000000009f800 (usable)

Enter <return> for next page or q<return> to go back to command line.
```

Turn paging off for this session:

```
Infoblox > set lines 0
Number of scroll lines set to 0.
```

Set a permanent line number:

```
Infoblox > set lines permanent 24
Number of scroll lines set to 24.
```

## set log\_txn\_id

The **set log\_txn\_id** command enables or disables the display of DHCP transaction IDs in syslog messages. By default, DHCP transaction ID logging is enabled. When you enable DHCP transaction ID logging, the appliance displays transaction IDs for the following packets:

- DHCPDISCOVER
- DHCPREQUEST
- DHCPRELEASE

- DHCPDECLINE
- DHCPINFORM

In Grid Manager, the transaction IDs are appended to the end of the corresponding syslog messages with a prefix of "TransID." You can view this information in the **Administrator** tab -> **Logs** tab -> **Syslog** tab of Grid Manager. When you enable this feature, you must restart DHCP service for the feature to take effect. When you disable this feature, you must perform a force restart services for the change to take effect. Use the [show log\\_txn\\_id](#) to display the current status of DHCP transaction ID logging.

## Syntax

```
set log_txn_id (ON|OFF)
```

Argument	Description
ON	Enables DHCP transaction ID logging on an appliance.
OFF	Disables DHCP transaction ID logging on an appliance.

## Example

### Enable DHCP transaction ID logging on an appliance

```
Infoblox > set log_txn_id ON
```

```
DHCP Transaction id logging turned ON
```

```
DHCP force restart services is required in order for the changed value to take effect
```

### Disable DHCP transaction ID logging on an appliance

```
Infoblox > set log_txn_id OFF
```

```
DHCP Transaction id logging turned OFF
```

```
DHCP force restart services is required in order for the changed value to take effect
```

### Sample syslog messages in the Syslog tab of Grid Manager:

#### When DHCP transaction ID logging is on:

```
2013-03-25T09:39:41+00:00 daemon (none) dhcpd[14434]: info DHCPINFORM from 10.0.0.199
```

```
via 10.120.20.182 TransID 78563412: not authoritative for subnet 10.0.0.0
```

```
2013-03-25T09:39:36+00:00 daemon (none) dhcpd[14434]: info DHCPDISCOVER
```



```
from  
cc:bb:cc:dd:ee:ff via 10.120.20.182 TransID 78563412
```

#### When DHCP transaction ID logging is off:

```
2013-03-25T09:39:39+00:00 daemon (none) dhcpd[14434]: info DHCPREQUEST  
for 10.0.0.199  
from cc:bb:cc:dd:ee:ff (dhcp-10-0-0-199) via 10.120.20.182
```

## set lom

The **set lom** command configures the LOM (Lights Out Management) settings for the IPMI interface. To view the current network settings for the IPMI interface, use the [show lom](#) command.

## Syntax

```
set lom
```

This command has no arguments.

## Example

```
Infoblox > set lom  
Enter LOM IP address: 10.1.1.22  
Enter LOM netmask: 255.255.255.0  
Enter gateway address [Default: 10.34.10.1]:  
LOM network settings:  
IP address: 10.34.10.42  
Netmask:255.255.255.0  
Gateway address: 10.34.10.1  
Is this correct? (y or n): y  
Are you sure? (y or n): y
```

## set max\_recursion\_depth

The **set max\_recursion\_depth** command sets the limit on the allowed number of levels of recursion named.

## Syntax

```
set max_recursion_depth <value>
```

Argument	Description
value	The depth value in the range 1 - 100. Default depth value is 7.

## Examples

```
Infoblox > set max_recursion_depth
```

## set max\_recursion\_queries

The `set max_recursion_queries` command sets the limit on the number of queries sent before terminating a recursive query.

## Syntax

```
set max_recursion_queries <value>
```

Argument	Description
value	The queries value in the range 1 - 1000. Default depth value is 150.

## Examples

```
Infoblox > set max_recursion_queries
```

## set membership

The `set membership` command specifies a Grid for the NIOS appliance. Use this command when the network address has been set (see [set network](#)) and you want to put the appliance in a Grid. You can join an IPv4 appliance to an IPv4-only or a dual mode Grid and an IPv6 appliance to an IPv6-only or a dual mode Grid. If the IP address is acceptable to the Grid Master, use this command to join the Grid. You can specify either an IPv4 or an IPv6 address of the Grid Master. If you need to re-address the appliance, use the [set network](#) command.

### Note

When you join a dual mode Grid member to a dual mode Grid, you can enter IPv4 address of the Grid Master if the Grid communication protocol for the Grid member is set as IPv4 and you can enter IPv6 address of the Grid Master if the Grid communication protocol for the Grid member is set as IPv6. For information about setting the communication protocol for a dual mode appliance, refer to the *Infoblox Administrator Guide*.

## Syntax

```
set membership
```

This command has no arguments.

## Example

```
Infoblox > set membership
Join status: No previous attempt to join a Grid.
Enter new Grid Master VIP: 10.1.1.22
Enter Grid Name [Default Infoblox]: DaveyJones
Enter Grid Shared Secret: L0ck37
Join Grid as member with attributes:
Join Grid Master VIP: 10.1.1.22
Grid Name: DaveyJones
Grid Shared Secret: L0ck37
WARNING: Joining a Grid will replace all the data on this node!
Is this correct? (y or n): y
Are you sure? (y or n): y
```

## set mgm attached

The set mgm attached command forces a Grid to attach to a Master Grid. Use this command only if a Grid is in the Attached state on the Multi-Grid Manager and Detached on the Grid Manager. This command recovers the Grid status when it is out of sync with the Grid status on the Multi-Grid Manager.

## Syntax

```
set mgm attached [MGM IP Address] [Port Number]
```

Argument	Description
MGM IP Address	IP address of the Master Grid
Port Number	Port number of the Master Grid

## Example

The following example uses the set mgm attached command.

```
Console connect [@ Grid IP address]
Infoblox > set maintenancemode
Maintenance Mode > set mgm attached [MGM IP address] [Port Number]
```

```
This command will force the Grid to get attached.
```

```
Are you sure you want to continue? (y or n): y
```

## set mld\_version\_1

The `set mld_version_1` command sets the IPv6 MLD (Multicast Listener Discovery) protocol to version 1, as described in *RFC 2710, Multicast Listener Discovery for IPv6*. MLD enables the appliance to detect multicast listeners on its directly attached links and discover which multicast addresses are of interest.

The appliance runs MLD version 2, as described in *RFC 3810, Multicast Listener Discovery Version 2 for IPv6*, by default. MLD version 2 is interoperable with version 1.

## Syntax

```
set mld_version_1
```

This command has no arguments.

## Example

```
Infoblox > set mld_version_1
Current MLD version: 2
Set Multicast Listener Discovery Version 1? (y or n): y New MLD Settings:
Use MLD version 1: Yes
Is this correct? (y or n): y MLD version: 1 is saved to database. MLD
version is set for IPv6.
```

## set monitor dns

The `set monitor dns` command enables network monitoring for DNS. Once enabled, you can do the following:

- View the average latency of authoritative and non-authoritative replies to DNS queries in 1, 5, 15, and 60 minute time intervals. Use the [show monitor](#) command to view the DNS network data.
- Monitor invalid DNS responses from UDP port 53. Use the [show monitor dns alert status](#) command to view the DNS alert status.

This command is useful when troubleshooting DNS and network issues.

### Note

This command is not supported for IPv6 in NIOS 7.0 and later releases. When you enable DNS network monitoring, there is a significant impact on DNS query performance.

## Syntax

```
set monitor dns {on | off}
```

Argument	Description
<code>on</code>	Enables network monitoring for DNS.
<code>off</code>	Disables network monitoring for DNS.

## Examples

### Turn on DNS network monitoring

```
Infoblox > set monitor dns on
Turning On DNS Network Monitoring...
```

### Turn off DNS network monitoring

```
Infoblox > set monitor dns off
Turning Off DNS Network Monitoring...
```

## set monitor dns alert

The `set monitor dns alert` commands enable DNS alert monitoring and set the thresholds for invalid DNS responses. After you enable DNS alert monitoring, the appliance monitors the UDP traffic on port 53 for recursive DNS queries, and then reports invalid DNS responses on UDP ports that are not open and with mismatched TXIDs. You must enable DNS network monitoring when you enable DNS alert monitoring. For information, see the [set monitor dns](#) command.

You can also configure the thresholds for invalid DNS responses. When the number of invalid responses exceeds the thresholds, the appliance logs the event and sends SNMP traps and notifications, if previously enabled. The default thresholds for both invalid ports and invalid TXIDs are 50%. You can configure the thresholds either as absolute packet counts or as percentages of the total traffic during a one minute time interval.

This command is useful for monitoring possible cache poisoning. Use the [show monitor dns alert status](#) command to view invalid port and invalid TXID data.



#### Note

This command is not supported for IPv6 in NIOS 7.0 and later releases.

## Syntax

```
set monitor dns alert {on | off}
set monitor dns alert modify {port | txid} over threshold_value {packets | percent}
```

Argument	Description
<code>on</code>	Enables DNS alert monitoring.
<code>off</code>	Disables DNS alert monitoring.
<code>modify</code> <code>port</code> <code>txid</code> <code>threshold_value</code> <code>packets</code> <code>percent</code>	Sets the thresholds for invalid DNS responses Enter <code>port</code> to set the threshold for invalid ports. Enter <code>txid</code> to set the threshold for invalid TXIDs. Enter the number of packets or percentage for the threshold. Enter <code>packets</code> if you want to set the threshold as a total packet count. Enter <b>percentage</b> if you want to set the threshold as a percentage of the total traffic. For a percentage-based threshold, the appliance does not generate a threshold crossing event if the traffic level is less than 100 packets per minute.

## Examples

### Turning on and off DNS alert monitoring

```
Infoblox > set monitor dns alert on
Infoblox > set monitor dns alert off
```

### Triggering a DNS alert when the percentage of invalid DNS responses on UDP ports exceeds 70% per minute

```
Infoblox > set monitor dns alert modify port over 70 percent
```

### Triggering a DNS alert when the total packet count of invalid DNS responses with mismatched TXIDs is over 100 packets per minute

```
Infoblox > set monitor dns alert modify txid over 100 packets
```

## set ms\_dns\_reports\_sync\_interval

You can use the `set ms_dns_reports_sync_interval` command to specify the time interval at which the DNS reporting data from the Microsoft server is synchronized with the NIOS appliance.

## Syntax

```
set ms_dns_reports_sync_interval <MS Server IP address> <seconds>
```

Argument	Description
<MS Server IP address>	Specify the IP address of the Microsoft server.
<seconds>	Specify the time interval in seconds at which the DNS reporting data from the Microsoft server is synchronized with the NIOS appliance. The default synchronization interval is 15 seconds.

## Example

```

Infoblox > set ms_dns_reports_sync_interval 10.102.30.2 14
Current DNS reports sync interval is 15 second(s).
The DNS reports sync interval will be changed to 14 second(s).
Is this correct? (y or n): y
The DNS reports sync interval has been changed to 14 second(s).

```

## set ms\_sticky\_ip

The `set ms_sticky_ip` command enables/disables `ms_sticky_ip`.

## Syntax

```
set ms_sticky_ip [on|off]
```

Argument	Description
on	Enables <code>ms_sticky_ip</code> .
off	Disables <code>ms_sticky_ip</code> .

## Examples

```
Infoblox > set ms_sticky_ip
```

## set named\_recv\_sock\_buf\_size

You can use the `set named_recv_sock_buf_size` command to tune the BIND receive socket buffer memory to a maximum of 8 MB. The DNS receive socket buffer holds BIND packets that are queued on the UDP (User Datagram Protocol) port from the NIC (Network Interface Controller). This command is useful when you want to increase the BIND receive buffer size to accommodate occasional burst traffic and high volume DNS recursive queries. Note that the same buffer is also used for updates and non-recursive queries. Use the [show named\\_recv\\_sock\\_buf\\_size](#) to view the current buffer size.

 **Note**

Ensure that you use this command only when you are dealing with burst traffic situations in high volume deployments.

## Syntax

```
set named_recv_sock_buf_size {N}
```

Argument	Description
N	The number of kilobytes to which you want to set the BIND receive socket buffer size. The minimum is 120 kilobytes and the maximum is 8192. The default is 1536.

## Example

Set the BIND receive socket buffer size to 5000 KB

```
Infoblox > set named_recv_sock_buf_size 5000
```

```
Infoblox >
```

## set named\_tcp\_clients\_limit

You can use the `set named_tcp_clients_limit` command to set the maximum number of simultaneous DNS clients that can be handled with TCP connections. It does not account for UDP connections.

## Syntax

```
set named_tcp_clients_limit <number of TCP clients>
```

Argument	Description
Number of TCP clients	Maximum number of simultaneous DNS clients that can be handled with TCP connections. The number must be between 200 and 25000. The default value is 1000.

## Example

```
Infoblox > set named_tcp_clients_limit 2500
```



## set network

The `set network` command specifies an address for a NIOS appliance so that it can join a network, with the option of joining a Grid. You can configure either IPv4 address, IPv6 address, or both for a NIOS appliance. If the appliance is configured with an IPv6 address, it can join a Grid using the IPv6 address of the Grid Master. Use the `show network` command to view the network settings.

`set network` supports configuration of both IPv4 and IPv6 interface addresses.

## Syntax

```
set network
```

This command has no arguments.

## Example

### Specifying an IPv4 address

```
Infoblox > set network
```

```
NOTICE: All HA configuration is performed from the GUI. This interface is used only to
```

```
configure a standalone node or to join a grid.
```

```
Enter IP address:10.35.1.104
```

```
Enter netmask [Default: 255.255.255.0]:
```

```
Enter gateway address [Default: 10.35.0.1]:
```

```
Enter VLAN tag [Default: Untagged]:
```

```
Enter DSCP value [Default: Inherited from Grid: 0]:
```

```
NOTICE: Additional IPv6 interface can be configured only via GUI.
```

```
Configure IPv6 network settings? (y or n):n
```

```
Become grid member? (y or n): n
```

```
New Network Settings:
```

```
IPv4 address: 10.35.1.104
```

```
IPv4 Netmask: 255.255.255.0
```

```
IPv4 Gateway address: 10.35.0.1
```

```
IPv4 VLAN tag: Untagged
```

```
IPv4 DSCP Value: Inherited from Grid: 0
```

```
Old IPv4 Network Settings:
```

```
IPv4 address: 192.168.1.2
IPv4 Netmask: 255.255.255.0
IPv4 Gateway address: 192.168.1.1
IPv4 VLAN tag: Untagged
IPv4 DSCP Value: Inherited from Grid: 0
```

## Specifying an IPv6 address

```
Infoblox > set network
NOTICE: All HA configuration is performed from the GUI. This interface is
used only to
configure a standalone node or to join a grid.
Enter IP address : 2620:10a:6000:2400::168
Enter IPv6 Prefix Length [Default: none]: 64
Enter IPv6 gateway [Default: none]: 2620:10a:6000:2400::1

Enter VLAN tag [Default: Untagged]:
Enter DSCP value [Default: Inherited from Grid: 0]
Configure IPv4 network settings? (y or n):n
Become grid member? (y or n): n

New Network Settings:

IPv6 address: 2620:10a:6000:2400::168/64
IPv6 Gateway address: 2620:10a:6000:2400::1
IPv6 VLAN tag: Untagged
IPv6 DSCP Value: Inherited from Grid: 0
```

## Specifying both IPv4 and IPv6 address

```
Infoblox > set network
NOTICE: All HA configuration is performed from the GUI. This interface is
used only to
configure a standalone node or to join a grid.
Enter IP address : 10.35.1.104
Enter netmask [Default: 255.255.255.0]: 255.255.0.0
Enter gateway address [Default: 10.35.0.1]:
Enter VLAN tag [Default: Untagged]:
Enter DSCP value [Default: Inherited from Grid: 0]
Configure IPv6 network settings? (y or n):y
```

```
Enter IPv6 address [Default: none]: 2620:10A:6000:2400::168
Enter IPv6 Prefix Length [Default: none]: 64
Enter IPv6 gateway [Default: none]: 2620:10A:6000:2400::1
Enter VLAN tag [Default: Untagged]:
Enter DSCP value [Default: 30]:
Become grid member? (y or n): n
```

#### New Network Settings:

```
IPv4 address: 10.35.1.104
IPv4 Netmask: 255.255.0.0
IPv4 Gateway address: 10.35.0.1
IPv4 VLAN tag: Untagged
IPv4 DSCP Value: Inherited from Grid: 0

IPv6 address: 2620:10a:6000:2400::168/64
IPv6 Gateway address: 2620:10a:6000:2400::1
IPv6 VLAN tag: Untagged
DSCP value: 30
```

#### Old IPv4 Network Settings:

```
IPv4 address: 192.168.1.2
IPv4 Netmask: 255.255.255.0
IPv4 Gateway address: 192.168.1.1
IPv4 VLAN tag: Untagged
IPv4 DSCP Value: Inherited from Grid: 0
```

#### Note

After you confirm your network settings, the Infoblox application automatically restarts.

After configuring the network settings, you cannot change the type of network connectivity of the appliance through CLI. For example, if the appliance is configured in IPv4-only mode, then you can change only the IPv4 interface settings through CLI. But the type of network connectivity for the appliance can be changed through GUI.

## set nogrid

The **set nogrid** command removes the specified member from the current Grid. Execute this command from the Grid member. This command is valid only on a member.

#### Note

Infoblox recommends that you use this command only in an emergency, such as when the network is down between the master and the member. Otherwise, you should configure the member to leave the Grid using the GUI on the Grid Master.

## Syntax

```
set nogrid
```

This command has no arguments.

## Example

```
Infoblox > set nogrid
```

```
The normal method to configure a node to leave a Grid is to use the GUI on the Grid Master. This method is only used for emergencies (e.g. network is down from the master to this node).
```

```
Is this such an emergency? y
```

```
The current node will become a standalone machine, with default values for Grid settings.
```

```
Are you sure? (y or n) y
```

```
The network settings have been updated.
```

## set nomastergrid

In a Multi-Grid environment, the **set nomastergrid** command enables a Grid to leave the current Master Grid. This command is valid only on the Multi-Grid Master.

## Syntax

```
set nomastergrid
```

This command has no arguments.

## Example

```
Infoblox > set nomastergrid
```

```
This grid is going to leave master grid
```

```
Are you sure? (y or n): y
```

```
Grid is not joined to a master grid. Exiting without making any change
```

## set nosafemode

The `set nosafemode` command disables safe mode on the NIOS appliance by re-enabling DNS and DHCP services. For more information, see [set safemode](#).

### Syntax

```
set nosafemode
```

This command has no arguments.

### Example

```
Infoblox > set nosafemode
```

## set oosp

The `set oosp off` command disables OOSP authentication service configuration and allow you to login without OOSP validation. It also terminate administrative sessions for users who are currently logged in.

### Syntax

```
set oosp off
```

There are no arguments for this command.

### Example

```
Infoblox > set oosp off
```

## set ospf

The `set ospf` command writes statistical information to syslog. This command provides informational data that can be helpful for diagnostic purposes. The statistical information is written (dumped) to syslog. When viewing the syslog file, lines with names such as `ospf statistics` are the OSPF statistical information. Use the [show ospf](#) command to view the OSPF settings.

Syslog level describes the types of messages that are sent to syslog. You can identify the syslog information by using the `level` option.



#### Note

To use this command, the NSQ software package must be installed.

## Syntax

```
set ospf log {/level}
```

Argument	Description
<i>level</i>	Writes OSPF statistics to syslog with a specific associated level. The supported log levels are: debugging, informational, notifications, warnings, errors, critical, alerts, and emergencies.

## Example

```
Infoblox > set ospf log alerts
```

## set overload\_bootp

The `set overload_bootp` command enables/disables overloading BOOTP packets.



### Note

To disable DHCP overload bootp packet control, DHCP force restart service is required in order for the changed value to take effect.

## Syntax

```
set overload_bootp
```

There are no arguments for this command.

## Example

```
Infoblox > set overload_bootp
```

## set pc\_domain add

You can use the `set pc_domain add` command to add category overrides for a specified domain. Use category 104 to add domains to the global proxy list. Queries to domains added to the global proxy list are proxied to the Multi-Service Proxy server when the **Enforce the global proxy list** parental control option is selected. For more information, see [Adding Subscriber Sites](#).

For information about the status of the command, see [show pc\\_domain](#). You can use the `set pc_domain delete` command to delete a specific domain.

## Syntax

```
set pc_domain add <domain> <category 1> <category 2 > <category 3>
```

### Note

- Category values must be integers.
- If the specified domain is cached in a DNS Cache Acceleration server and if you add the same domain to the global proxy list using the `set pc_domain add` command, the domain does not get proxied unless the TTL (Time To Live) setting of the domain in the DNS Cache Acceleration server has expired. For more information about TTL settings, see [Specifying Time To Live Settings](#).
- Domains added to category 104 using the `set pc_domain add` command are not proxied for members until the domains are synchronized with the members, which happens only after the expiry of the time set in the **Update Interval in hours** field in Subscriber Services Properties. By default this is 24 hours. For more information about the setting, see [Enabling Subscriber Parental Control](#).

Argument	Description
<code>domain</code>	Name of the domain for which the category override is executed.
<code>category 1</code>	Category 1 value.
<code>category 2</code>	Category 2 value.
<code>category 3</code>	Category 3 value.

## Examples

```
Infoblox > set pc_domain add bbc.com 10308 0 0
```

```
Successfully added category override for "bbc.com".
```

## set pc\_domain delete

You can use the `set pc_domain delete` command to delete all category override for the specified domain. For information about the add category override command, see [set pc\\_domain add](#).

## Syntax

```
set pc_domain delete <domain>
```

Argument	Description
domain	The name of the domain for which the category override will delete.

## Example

```
Infoblox > set pc_domain delete bbc.com
```

```
Successfully deleted category override for "bbc.com".
```

## set phonehome

The **set phonehome** command enables a Grid Master or an independent appliance to email reports monthly and after each upgrade to Infoblox Technical Support and other specified recipients.

The reports provide status and event information about the Grid or independent appliance and its services. The report is an XML document that includes the following information:

- The phone home feature version.
- The report type, such as periodic and test.
- The time of the report.
- The Infoblox Support ID that was assigned to the account.
- Information about the Grid, such as its NIOS version, name, VIP, Grid Master hostname, LAN IP, and the number of Grid members and appliances in the Grid.
- The upgrade history of the Grid.
- Information about each Grid member, such as the hostname, IP address, status, role (such as standalone, master), and if the member is an HA pair. If the member is a peer in a DHCP failover association, the report also includes the DHCP failover status.
- Hardware information, such as the hardware type, serial number, HA status, and uptime.
- Information about the interfaces, such as the interface name and IP addresses.
- Resource usage information, such as CPU and system temperature, and CPU, database, disk, and memory usage.

## Syntax

```
set phonehome {on | off}
```

Argument	Description
on	Enables the appliance to send status and event reports to specified recipients.
off	Disables the function to send reports.



## Examples

### Turning on the phone home feature

```
Infoblox > set phonehome on
```

### Turning off the phone home feature

```
Infoblox > set phonehome off
```

## set promote\_master

The `set promote_master` command specifies a NIOS appliance as the new Grid Master in the case of a Grid Master failure. The new Grid Master then alerts all the Grid members to redirect their traffic to it. If you have configured multi-site reporting cluster, you can modify the primary reporting site. For information about reporting clusters, refer to the *Infoblox NIOS Administrator Guide*.

You can do one of the following to promote a master candidate to a Grid Master:

- Immediately notify all Grid members about the promotion.
- Set a sequential notification to provide wait time for Grid members to join the new Grid Master. Staggering the restarts of Grid members can minimize DNS outages. The sequential order for Grid members to join the new Grid Master begins with the old Grid Master and then the Grid members in FQDN order. The default delay time is 120 seconds. You can configure the delay time from a minimum of 30 seconds up to 600 seconds.

For this command to be effective, you must have previously specified an appliance as the Grid Master candidate. Then when you lose the Grid Master, you can remotely (SSH) log in to the Grid Master candidate and execute this command.



### Note

When the previous Grid Master comes back on line, it automatically joins the Grid as a master candidate.

## Syntax

```
set promote_master
```

This command has no arguments.

## Examples

```
Infoblox > set promote_master
```

```
Do you want a delay between notification to Grid members? (y or n):
```

Enter `n` to promote the master candidate and send notifications to all Grid members immediately. The appliance displays the following:

This action will immediately promote master candidate to become the Grid Master. This feature is designed to be used primarily for disaster recovery.

Are you sure you want to do this? (y or n): y

The current member will become the Grid Master.

Are you really sure you want to do this? (y or n): y

Member promotion beginning on this member.

Enter y to promote the master candidate to the Grid Master immediately and specify the delay time for the Grid members to join the new Grid Master. The appliance displays the following:

Set delay time for notification to Grid members? [Default: 120s]: 200

This action will immediately promote master candidate to become the Grid Master. This feature is designed to be used primarily for disaster recovery.

Are you sure you want to do this? (y or n): y

The current member will become the Grid Master. The Grid members will be notified sequentially with a delay of 200 seconds.

Are you really sure you want to do this? (y or n): y

If you have configured multi-site reporting cluster, the appliance displays all the reporting sites in the order of priority you have configured. For example if you have configured the following reporting sites: site 4(priority 1), site 2 (priority 2), site 1(priority 3), and site 3 (priority 4)

```
Infoblox > set promote_master
```

Do you want a delay between notification to Grid members? (y or n): n

Primary reporting site candidates (in order of priority):

1. site4 (Existing primary reporting site)
2. site2
3. site1
4. site3

Please enter new primary reporting site (1-2) or 'c' to continue without changing primary reporting site: 5

The appliance displays the following error when you enter value incorrectly:

ERROR: Please enter a valid choice or 'c' to continue without changing the primary reporting site.

Please enter new primary reporting site (1-2) or 'c' to continue without changing primary reporting site: c

This action will immediately promote this member to become the grid master. This feature is designed to be used primarily for disaster recovery.

Are you sure you want to do this? (y or n): y The current member will become the grid master.

Are you really sure you want to do this? (y or n): y

```
Master promotion beginning on this member
```

```
Good Bye
```

To change the primary reporting site:

```
Infoblox > set promote_master
```

```
Do you want a delay between notification to Grid members? (y or n):n
```

```
Primary reporting site candidates (in order of priority):
```

1. site4 (Existing primary reporting site)
2. site2
3. site1
4. site3

```
Please enter new primary reporting site (1-4) or 'c' to continue without  
changing primary reporting site: 2
```

```
Are you sure you want to do this? (y or n): y The current member will  
become the grid master.
```

```
Are you really sure you want to do this? (y or n): y Master promotion  
beginning on this member
```

```
Good Bye
```

The new priority order of reporting sites will be:

```
site2 (Existing primary reporting site)
```

```
site4
```

```
site1
```

```
site3
```

## set gmc\_promotion

Use the `set gmc_promotion disable` command to disable the **Activate GMC Group Promotion Schedule** option. Note that, this command can be executed on Grid Master and Grid Master Candidate.

Use the `set gmc_promotion forced_end` command while Grid Master Candidate promotion is in progress, to unset the master promotion status. Note that, this command can be executed only on master Node.

## Syntax

```
set gmc_promotion <disable >< forced_end >
```

Argument	Description
<code>disable</code>	Disables the <b>Activate GMC Group Promotion Schedule</b> option.

Argument	Description
<code>forced_end</code>	Resets the Grid Master Candidate promotion flag to reconfigure Grid Master Candidate group promotion.

## Examples

### Disabling the Grid Master Candidate group promotion

```
Infoblox > set gmc_promotion disable
Feature is disabled.
```

### Resetting the flag to reconfigure Grid Master Candidate group promotion

```
Infoblox > set gmc_promotion forced_end
This command will not end GMC promotions but just resetting the flag for end-users to
reconfigure GMC-promotion again.
This command can be executed only on Master node.
```

## set prompt

Use the `set prompt` command to change the prompt to the host name, `user@host name`, host IP address, or `user@ host IP address`. Note that the prompt displayed in the command line interface (CLI) can be set only on the active Grid Master node. Once you execute the `set prompt` command, the prompt displayed for all Grid members is set accordingly and you can see the prompt when you log in to the CLI for each Grid member.

## Syntax

```
set prompt {hostname / user@hostname / ip / user@ip / default}
```

Argument	Description
<code>hostname</code>	Sets the prompt to the host name of the computer from which you access the appliance.
<code>user@hostname</code>	Sets the prompt to the user name@ the host name of the computer from which you access the appliance.
<code>ip</code>	Sets the prompt to the IP address of the host.
<code>user@ip</code>	Sets the prompt to the user name@ the IP address of the host.

Argument	Description
default	Sets the prompt to "Infoblox >".

## Example

```
Infoblox > set prompt user@hostname
admin@infoblox >
admin@infoblox > set prompt user@ip
admin@172.31.1.254 >
```

## set regenerate\_anycast\_password

The `set regenerate_anycast_password` command regenerates eight-character alphanumeric password for the anycast service and saves it to the NIOS database. The password can then be used across all anycast configuration files (`ospf.conf`, `bgp.conf`, `bfd.conf`) with the following CLI commands:

- `show ospf`
- `show bgp`
- `show ipv6_ospf`
- `show ipv6_bgp`
- `show bfd`

When these commands are executed with the configuration argument, the values of the password and enable password fields are shown encrypted. Running `set regenerate_anycast_password` restarts the anycast service on the Grid members that have the anycast service enabled. For information on OSPF commands, see [show ospf](#) and [show ipv6\\_ospf](#). For information on BGP commands, see [show bgp](#) and [show ipv6\\_bgp](#). For information on bfd command, see [show bfd](#).

### Note

Only superusers can use `set regenerate_anycast_password` in maintenance mode.

## Syntax

```
set regenerate_anycast_password
```

This command has no arguments.

## Example

```
Maintenance Mode > set regenerate_anycast_password
Resetting the password will restart the anycast service.
Do you want to proceed? (y or n):y
```

```
Are you sure? (y or n):y
Anycast password generation is successful.
```

## set remote\_console

The `set remote_console` command enables and disables access to the NIOS appliance using a remote console. Use the `show remote_console` command to view the remote console settings.

### Note

Infoblox recommends that you close any port that is not being used, for security reasons. An open, unused port offers the potential for unwanted access to your network.

## Syntax

```
set remote_console
```

This command has no arguments.

## Example

```
Infoblox > set remote_console
Enable remote console access (Grid level)? (y or n): y
New remote console access settings:
Remote console access enabled: Yes
Is this correct? (y or n): y
```

## set reporting\_cert

In a Grid with a reporting server, you can use the `set reporting_cert` command to generate a new set of SSL certificates on all forwarders and the indexer. You can use this command only on the Grid Master.

## Syntax

```
set reporting_cert
```

This command has no arguments.

## Example

```
Infoblox > set reporting_cert
Generate new reporting certificate? (y or n): y
Reporting certificates generated.
```

## set reporting\_user\_capabilities

The `set reporting_user_capabilities` command allows you to configure the `delete` permission on reporting data to a local admin user who has superuser permissions. If you enable the `set reporting_user_capabilities` command for a user, the user can use the `delete` command using the [Splunk API](#) or [reporting GUI](#) to delete selected events.

To see the list of users configured with the reporting delete permission, see [show reporting\\_user\\_capabilities](#).

### Notes

- This command is supported only on the Grid master.
- You cannot retrieve the data once it is deleted.
- The deleted data cannot be visualized and does not reduce any disk space.
- Frequent deletion of data may affect the search performance.

## Syntax

```
set reporting_user_capabilities [enable|disable] <super-user>
```

Argument	Description
enable	Enables the reporting delete capability
disable	Disables the reporting delete capability

## Example

```
Infoblox > set reporting_user_capabilities enable user1
```

```
1. Delete reporting indexed data
```

```
Select capability (1) or q to quit: 1
```

```
The reporting Delete capability has been enabled for user user1.
```

```
Infoblox > set reporting_user_capabilities disable user1
```

```
1. Delete reporting indexed data
```

```
Select capability (1) or q to quit: 1
```

```
The reporting Delete capability has been disabled for user user1.
```

## set restart\_anycast\_with\_dns\_restart

The `set restart_anycast_with_dns_restart` command sets the DNS and anycast start and restart sequence to the appliance. This command stops the anycast service during a restart or stop of the DNS service and redirects the traffic present on the anycast IP address to another site. You can use this command only on Grid Master. To view information about the current status of the command, see [show restart\\_anycast\\_with\\_dns\\_restart](#).

### Syntax

```
set restart_anycast_with_dns_restart [ on | off ]
```

Argument	Description
on	Restarting the DNS service stops the anycast service followed by a DNS service stop. The anycast service restarts after the DNS service loads all the zones. In case the DNS service stops, anycast service also stops. The DNS service start triggers anycast service start, after the DNS service loads all the zones.
off	Restarting the DNS service will not trigger the restart of anycast service. In case of the DNS service stop, anycast service stops after approximately 30-60 seconds of DNS service stop.  In case of DNS service start: <ul style="list-style-type: none"><li>• If the DNS service takes more than approximately 30 seconds to load all the zones, anycast service starts at approximately 30 seconds after DNS service starts.</li><li>• If the DNS service loads immediately, anycast service starts immediately after DNS service starts.</li></ul>

### Example

The following is an example to restart anycast.

```
Infoblox > set restart_anycast_with_dns_restart on
```

## set revert\_grid

Use the `set revert_grid` command to revert to a version of software that was running previously on a Grid or on an independent appliance or HA pair. Be aware that when you revert to this software, any configurations made to the currently running software are lost. You can back up the current data before you revert so that you can later determine what configuration changes are missing.

### Syntax

```
set revert_grid
```

This command has no arguments.



## Example

```
Infoblox > set revert_grid
```

## set rpz\_recursive\_only

Use the `set rpz_recursive_only <view_name> [<zone_name>]` command to use NIOS RPZ zones instead of local RPZ zones to block records with private IP addresses from being queried by external users. This command is available only on the Grid Master.

If you do not specify an RPZ zone name after the DNS view name, all RPZ zones that belong to the specified DNS view are used to block records. If you specify an RPZ zone name, only that zone is used to block records.



Restart the DNS service on the member after running the command.

## Syntax

```
set rpz_recursive_only <view_name> [<zone_name>] <none | yes | no>
```

Argument	Description
view_name	DNS view to which the RPZ zones belong.
zone_name	NIOS zone name that must be used to block records with private IP addresses. If you do not specify an RPZ zone, all zones that belong to the DNS view are considered.
none	Uses the existing setting for the command. For example, if the <code>set rpz_recursive_only view_name</code> command was set to <code>yes</code> , specifying <code>none</code> will consider the command enabled because the earlier setting was set to <code>yes</code> .
yes	Enables the command.
no	Disables the command.

## Example

```
Infoblox > set rpz_recursive_only default rpz1.com yes
Restart the DNS service in order for changes to take effect.
```

## set safemode

The `set safemode` command disables DNS and DHCP services. Use this command to troubleshoot a NIOS appliance with unreliable services.

This command restarts all the services, including DNS and DHCP. DNS and DHCP remain active only long enough to write `named.conf` and `dhcp.conf` files. These services then shut down. All other services remain functional. This

allows you to review the `named.conf` and `dhcp.conf` files to determine and alleviate the cause of the appliance distress.

Once you have determined the problem, you can reinstate DNS and DHCP services using the [set nosafemode](#) command.

## Syntax

```
set safemode
```

This command has no arguments.

## Example

```
Infoblox > set safemode
```

## set scheduled

Use the `set scheduled` command to specify the number of times per hour the appliance checks if the services need a restart when the task scheduling feature is enabled. You must manually restart services or schedule a restart of services for the scheduled change to take effect.

You can set the value from 0 to 60, and the default value is 60. When you set the value to 0, the appliance turns off the restart feature.

Use the [show scheduled](#) to view the number of times per hour the appliance checks whether a restart of services is required.

## Syntax

```
set scheduled task restarts [0-60]
```

Argument	Description
0-60	The number of times per hour the appliance checks if the services need a restart when the task scheduling feature is enabled. You can enter any number from 0 to 60. The default is 60. A value of 0 turns off the restart feature.

## Example

Enter the following command to enable the appliance to check 10 times per hour whether the services need a restart:

```
Infoblox > set scheduled task restarts 10
```

The appliance checks 10 times per hour if the services must be restarted, which is every six minutes of the hour. For example, if you enter the command at 3:15 p.m., the appliance checks if the services must be restarted every six minutes starting at the hour (3:00 p.m.). Therefore, the next checks are at 3:18, 3:24, 3:30, 3:36, 3:42, 3:48, 3:54, and 4:00 p.m.

## set security

The `set security` command allows you to specify IP or network addresses that can access the appliance through the GUI. The appliance denies access to addresses that are not specified. Use the `show security` command to view the security settings.

### Syntax

```
set security
```

This command has no arguments.

### Example

In the following example, security is enabled to restrict access to the NIOS appliance (through the GUI) to the IP address range 10.1.1.1:

```
Infoblox > set security
Enable security? (y or n): y
Enter access IP range: 10.1.1.1
Enter access netmask (Default: 255.255.255.0): 255.255.255.0
New security settings:
Security enabled: Yes
IP range: 10.1.1.1
Is this correct? (y or n): y
Do you wish to enter additional access range? (y or n): n
```

## set session\_timeout

Use the `set session_timeout` command to specify how long a session remains open when there is no user activity. Use the `show session_timeout` command to view the `session_timeout` setting.

### Syntax

```
set session_timeout
```

This command has no arguments.

## Example

```
Infoblox > set session_timeout
```

```
Current GUI/CLI timeout is 60000 seconds (16:40:00)
```

```
WARNING: Changing the session timeout will cause GUI users to be logged out.
```

```
New GUI/CLI session timeout (in seconds, 0 to abort)? 90000
```

## set smartnic monitor-mode

The `set smartnic monitor-mode` command enables and disables monitor mode for the Threat Protection service. This is disabled by default. When monitor mode is enabled, the appliance logs DNS packets (instead of dropping them) that would have been blocked by threat protection rules. This information is recorded in the audit log. Note that you can enable or disable monitor mode only for individual members. You cannot set this configuration at the Grid level.

To view whether monitor mode is enabled or disabled for the Threat Protection service, see [show smartnic](#).

## Syntax

```
set smartnic monitor-mode {on|off}
```

Argument	Description
on	Enables monitor mode for the Threat Protection service.
off	Disables monitor mode for the Threat Protection service.

## Examples

### Enable debugging

```
Infoblox > set smartnic monitor-mode on
```

### Disable debugging

```
Infoblox > set smartnic monitor-mode off
```

## set snmptrap

The `set snmptrap` command sends SNMP traps to the trap receiver (SNMP trap server IP address) you specify.

You can use the optional `v3` command to generate SNMPv3 traps. For information about SNMP, see [Monitoring with SNMP](#).

Use the `show snmp` command to get information about SNMP objects.



### Note

Use the `set snmptrap` command only for test purposes because it uses 0 as the value of the `msgAuthoritativeEngineBoots` and `msgAuthoritativeEngineTime` variables and this may cause the trap receiver to drop traps.

## Syntax

```
set snmptrap variable <name of an SNMP variable, in dotted or symbolic format>
address <address of the trap receiver> [v3] [snmpuser][all] [ibName <value> ]
[ibTrapSeverity <value>] [ibObjectName <value>] [ibProbableCause
<value>] [ibSubSystemName <value>] [ibCurThresholdValue <value>]
[ibThresholdHigh <value>] [ibThresholdLow <value>] [ibPreviousState <value>]
[ibCurrentState <value>] [ibTrapDesc <value>]
```

Argument	Description
<name of an SNMP variable>	Name or OID (object ID) of the SNMP object. For example, you can enter sysName.0 or .1.3.6.1.4.1.2021.11.53.0.
<address of the trap receiver>	IPv4 or IPv6 address of the management system that receives SNMP traps.
<snmpuser>	User name of the SNMPv3 user account. This is optional. If you do not provide a user name, the appliance uses the first SNMPv3 user on the list.
ibName <value>	IP address of the appliance to which the trap must be sent.
ibTrapSeverity <value>	Severity of the trap. For more information, see <a href="#">Trap Severity (OID 3.1.1.1.2.2.0)</a> .
ibObjectName <value>	Name of the object for which the trap is generated.

Argument	Description
<code>ibProbableCause &lt;value&gt;</code>	Value that provides information about events such as hardware, software for process failures that trigger SNMP traps.
<code>ibSubSystemName &lt;value&gt;</code>	Value that provides information about the subsystems that trigger the traps.
<code>ibCurThresholdValue &lt;value&gt;</code>	Current value of the threshold counter.
<code>ibThresholdHigh &lt;value&gt;</code>	For CPU usage, this is the trigger value of the SNMP trap. For DHCP address usage, this is the value of the high watermark.
<code>ibThresholdLow &lt;value&gt;</code>	For CPU usage, this is the reset value of the SNMP trap. For DHCP address usage, this is the value for the low watermark.
<code>ibPreviousState &lt;value&gt;</code>	Previous state of the appliance.
<code>ibCurrentState &lt;value&gt;</code>	Current state of the appliance.
<code>ibTrapDesc &lt;value&gt;</code>	Description of the trap.

For more information about the argument descriptions, see [Types of Traps \(OID 3.1.1.1.1\)](#).

## Examples

### Sending SNMP Traps to a Specific Trap Receiver

Enter the following on the appliance:

```
Infoblox > set snmptrap variable sysName.0 address 10.0.0.11
```

The appliance sends the following acknowledgement to the trap receiver:

```
2011-02-23 23:02:51 10.0.0.11 [UDP: [10.0.0.11]:35597->[10.0.0.11]]:
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (42) 0:00:00.42
SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-MIB::sysName.0
2011-02-23 23:02:53 10.0.0.11 [UDP: [10.0.0.11]:52367->[10.0.0.11]]:
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (42) 0:00:00.42
SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-MIB::sysName.0
```

## Sending SNMP Traps Using SNMPv3

Enter the following on the appliance:

```
Infoblox > set snmptrap variable sysName.0 localhost v3 SNMPv3User1
```

The appliance sends the following acknowledgement to the trap receiver:

```
2011-02-07 01:08:19 localhost [UDP: [127.0.0.1]:41884->[127.0.0.1]]:  
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (42) 0:00:00.42  
SNMPv2-MIB::snmpTrapOID.0 = OID: DISMAN-EVENT-MIB::sysName.0
```

Enter the following on the appliance:

```
Infoblox > set snmptrap variable sysName.0 localhost v3 SNMPv3User1
```

The appliance sends the following to the trap receiver in the event of a process failure:

```
SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.7779.3.1.1.1.1.2.0  
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.1.0 = STRING: "192.168.1.2"  
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.2.0 = INTEGER: 5  
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.5.0 = STRING: "named"  
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.4.0 = INTEGER: 20  
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.11.0 = STRING: "A named daemon  
monitoring  
failure has occurred."
```

## Sending SNMP Traps When the Primary Disk is Full

Enter the following on an appliance of type disk:

```
Infoblox > set snmptrap variable 1.3.6.1.4.1.7779.3.1.1.1.1.3 address  
10.35.5.243 ibNodeName "infoblox.localdomain" ibObjectName Disk ibTrapSeverity 3  
ibProbableCause 1001 ibTrapDesc "'Primary drive is full.'"
```

The appliance sends the following to the trap receiver in the event of the primary disk being full:

```
2020-06-30 12:36:25 member0.g.infoblox.com [UDP: [10.35.5.243]:44807-  
>[10.35.5.243]:162]:  
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (10123392) 1 day, 4:07:13.92  
SNMPv2-MIB::snmpTrapOID.0 = OID: IB-TRAP-MIB::ibThresholdCrossingEvent IB-TRAP-  
MIB::ibNodeName.0 = STRING: "infoblox.localdomain" IB-TRAP-  
MIB::ibTrapSeverity.0 = INTEGER: minor(3) IB-TRAP-MIB::ibObjectName.0 =  
STRING: Disk IB-TRAP-MIB::ibCurThresholdValue.0 = INTEGER: 100 IB-TRAP-
```

```
MIB::ibThresholdHigh.0 = INTEGER: 90    IB-TRAP-MIB::ibThresholdLow.0 =  
INTEGER: 0    IB-TRAP-MIB::ibTrapDesc.0 = STRING: Primary drive is full.
```

## set static\_route

Use the `set static_route` command to configure static routes on your appliance. You can create new IPv4 and IPv6 static routes, move static route to a desired position in the list of static routes, move static routes up and down in the list, or move static route at a specific position up or down in the list of static routes. You can also delete an existing static route, delete static routes at a specific position, or delete all existing static routes for IPv4, IPv6 or both protocols. Use the [show static\\_routes](#) command to view the current configuration of static routes on your appliance.

## Syntax

```
set static_route add network/cidr gateway [position]  
set static_route move network/cidr gateway new-position  
set static_route move network/cidr gateway (up|down)  
set static_route move (v4|v6) old-position new-position  
set static_route move (v4|v6) old-position (up|down)  
set static_route delete network/cidr gateway  
set static_route delete (v4|v6) position  
set static_route delete (all|v4|v6)
```

Argument	Description
<code>add network/cidr gateway [position]</code>	Creates a new static route and optionally specifies its position in the list of static routes.
<code>move network/cidr gateway new-position</code>	Moves a static route to a new position in the list of static routes.
<code>move network/cidr gateway [up down]</code>	Moves a static route up or down in the list of static routes.
<code>move (v4   v6) [old-position] [new-position]</code>	Moves an IPv4 or IPv6 static route from its existing position to a new position in the list of static routes.
<code>move (v4   v6) [old-position] (up down)</code>	Moves an IPv4 or IPv6 static route up or down from its current position in the list of static routes.
<code>delete network/cidr gateway</code>	Deletes an existing static route.
<code>delete (v4   v6) [position]</code>	Deletes an IPv4 or IPv6 static route, at a specific position, from the list of static routes.
<code>delete (all   v4   v6)</code>	Deletes all static routes, all IPv4 static routes, or all IPv6 static routes.



## Examples

### Adding a new static route

```
Infoblox > set static_route add 10.6.112.0/24 10.6.42.1  
Infoblox > set static_route add 2001:1234:5678::/112 2001:1234::42
```

### Moving a static route to a new position in the list of static routes

```
Infoblox > set static_route move 10.6.112.0/24 10.6.42.1 5  
Infoblox > set static_route move 2001:1234:5678::/112 2001:1234::42 6
```

### Moving a static route up or down in the list of static routes

```
Infoblox > set static_route move 10.6.112.0/24 10.6.42.1 up  
Infoblox > set static_route move 10.6.112.0/24 10.6.42.1 down  
Infoblox > set static_route move 2001:1234:5678::/112 2001:1234::42 up  
Infoblox > set static_route move 2001:1234:5678::/112 2001:1234::42 down
```

### Moving an IPv4 or IPv6 static route from its current position to a new position in the list of static routes

```
Infoblox > set static_route move v4 1 12  
Infoblox > set static_route move v6 1 12
```

### Moving an IPv4 or IPv6 static route up or down from its current position in the list of static routes

```
Infoblox > set static_route move v4 12 up  
Infoblox > set static_route move v6 12 up  
Infoblox > set static_route move v4 1 down  
Infoblox > set static_route move v6 1 down
```

### Deleting an existing static route

```
Infoblox > set static_route delete 10.6.112.0/24 10.6.42.1  
Infoblox > set static_route delete 2001:1234:5678::/112 2001:1234::42
```

## Deleting an IPv4 or IPv6 static route at a specific position in the list of static routes

```
Infoblox > set static_route delete v4 1
```

```
Infoblox > set static_route delete v6 2
```

## Deleting all IPv4 static routes

```
Infoblox > set static_route delete v4
```

## Deleting all IPv6 static routes

```
Infoblox > set static_route delete v6
```

## Deleting all IPv4 and IPv6 static routes

```
Infoblox > set static_route delete all
```

## set subscriber\_secure\_data add

If you have configured Infoblox Subscriber Services, you can use the `set subscriber_secure_data add` command to add a specific subscriber record to the subscriber database. For information about Infoblox Subscriber Services, see [Infoblox Subscriber Services](#). You can use `set subscriber_secure_data delete` command to delete a subscriber record.

## Syntax

```
set subscriber_secure_data add <ip_addr> <prefix> <local_id> <ip_space_desc>  
<data_string> [flags]
```

Argument	Description
<ip_addr>	Specify the IPv4 or IPv6 address of the subscriber.
<prefix>	Specify the prefix length.
<local_id>	Specify the Local ID, which is the MAC address of the subscriber device. To indicate a record without a Local ID, enter N/A.
<ip_space_desc>	Specify the IP space discriminator. To indicate a record without an IP space discriminator, enter N/A.
<data_string>	The cache data string is a concatenation of "AVPTAG:AVP-NAME=Value" separated by ';'. Example: AVPTAG:AVP-NAME=10.10.10.10;AVPTAG:AVP-NAME=10.10.10.10

Argument	Description
[flags]	This is optional. You can specify the following: S - For the non-expiring static record. B - If a blacklist or whitelist domains are configured for the subscriber.

The AVPTAG can include the following tags:



#### Note

The following tags are for the AVPs available in the list of predefined AVPs. You can also add tags for user-defined AVPs.

- **SUB**: Indicates the subscriber ID AVP, that is configured in the *Subscriber Services Properties* editor of the Grid. Example: `SUB:IMSI=602030100000057`.
- **NAS**: Indicates the NAS contextual information AVP, that is configured in the *Subscriber Services Properties* editor of the Grid. Example: `NAS:NAS-PORT=168496141`.
- **IPS**: Indicates the IP space discriminator AVP, that is configured in the *Subscriber Services Properties* editor of the Grid. Example: `IPS:NAS-IPv6-Address=2620:010a:6000:22f2::006e`.
- **PCP**: Parental-Control-Policy. It is a 128-bit value provided by the Service Providers. Example: `PCP:Parental-Control-Policy=400000`.
- **DCP**: Dynamic-Category-Policy. It is an indication to resolve domains categorized as dynamic. If Dynamic AVP is set, then all domains categorized as dynamic are forwarded to Infoblox Harmony. If Dynamic AVP is not set, then the domains categorized as dynamic would either be blocked or allowed depending on the parental control policy. Example: `DCP:Dynamic-Category-Policy=0`.
- **SSP**: Subscriber-Secure-Policy. It is a 32-bit value provided by the Service Providers. Example: `SSP:Subscriber-Secure-Policy=5F`.
- **EXP**: The date and time when the profile expires. Example: `EXP:Expire-Profile=Mon May 29 10\ \:23\ \:56 2017`.
- **AN0**: Indicates the AVP configured in ancillary position 0 in the ancillary list, that is configured in the *Subscriber Services Properties* editor of the Grid. Example: `AN0:Class=0x010x000x120x000x12`.
- **AN1**: Indicates the AVP configured in ancillary position 1 in the ancillary list, that is configured in the *Subscriber Services Properties* editor of the Grid. Example: `AN1:IMEI=8635800299072601`.
- **AN2**: Indicates the AVP configured in ancillary position 2 in the ancillary list, that is configured in the *Subscriber Services Properties* editor of the Grid. Example: `AN2:NAS-Port-Identifier=portid-1232`.
- **AN3**: Indicates the AVP configured in ancillary position 3 in the ancillary list, that is configured in the *Subscriber Services Properties* editor of the Grid. Example: `AN3:User-Name=user1`.

---

**Note:** The ancillary fields are numbered as AN0, AN1, AN2 etc. You cannot add more than five ancillary AVPs.

---

- **ACS**: The accounting session ID. Example: `ACS:Acct-Session-Id=29de847acde415ab`.

## Example 1

```
set subscriber_secure_data add 2620:10a:6000:7814::50b 128 003048d5d928 N/A
"ACS:Acct-Session-Id=29de847acde415ab;LID:003048d5d928;IPS:NAS-IPv6-
Address=2540:010a:6000:22f2::006e;AN0:APN=corp1;SUB:IMSI=602030100000045;AN3
:User-Name=user1;AN2:NAS-Port-Identifier=portid-1232;NAS:NAS-
PORT=168496141;SSP:Subscriber-Secure-Policy=00000001;DYN:Dynamic-Category-
Policy=0;PCP:Parental-Control-Policy=400000;"
```

## Example 2

```
set subscriber_secure_data add 10.32.1.145 32 AABCC112233 corp1 "ACS:Acct-
Session-
Id=32de327aced215ab;SUB:IMSI=301030100000026;LID:AABCC112233;NAS:NAS-
PORT=168496141;PCP:Parental-Control-Policy=20001;EXP:Expire-Profile=Mon May
29 10\\:23\\:56 2017;DCP:Dynamic-Category-Policy=0;SSP:Subscriber-Secure-
Policy=5F;"
```

## Format of a Subscriber Record

The following is an example of a subscriber record in the subscriber cache:

```
10.32.1.145/32 | IPS:corp1 | ACS:Acct-Session-
Id=32de327aced215ab;SUB:IMSI=301030100000026;LID:003048d5d928;NAS:NAS-
PORT=168496141;PCP:Parental-Control-Policy=20001;EXP:Expire-Profile=Mon May
29 10\\:23\\:56 2017;DCP:Dynamic-Category-Policy=0;SSP:Subscriber-Secure-
Policy=5F | 2017-06-05 21:20:51
```

## set subscriber\_secure\_data bypass

You can use the `set subscriber_secure_data bypass` command to configure bypass subscriber service policies on all members of each site on the entire Grid, all members of the site or on a local member. To view information about the status of the command, see [show subscriber\\_secure\\_data bypass](#).

## Syntax

```
set subscriber_secure_data bypass <on | off > [grid | site][ "site_name" ]
```

Argument	Description
<code>on</code>	Enables subscriber service policies bypass.
<code>off</code>	Disables subscriber service policies bypass.
<code>grid</code>	Enable or disable the subscriber service policies bypass for all members of each site on the entire Grid.
<code>site</code>	Enable or disable the subscriber service policies bypass for all members of the site.

 **Note**

- The newly added member do not inherit the bypass settings configured for the site. You must re-execute the bypass on/off CLI command.
- For a member of the group, that is not upgraded, the bypass subscriber service policies are not configured, and the message "Failed to set Subscriber Secure Bypass state for the member dns.com: upgrade is in progress." is displayed.
- The member must be part of a site even if the Subscriber Services are not running in order to configure Bypass Subscriber Service Policies.
- When there is a DNS service restart, if a member is restarting and the `set subscriber_secure_data bypass` command is executed, then the Subscriber Secure Bypass state is not enabled for the member.
- You can execute the `set subscriber_secure_data bypass` command with `grid` and `site` arguments only on the Grid Master and not on members.

## Examples

```
Infoblox > set subscriber_secure_data bypass on site mobile1
```

```
Site: mobile1
```

```
Member: dns1.com
```

```
Subscriber Secure Bypass enabled for the member
```

```
dns1.com.
```

```
Member: dns2.com
```

```
Subscriber Secure Bypass enabled for the
```

```
member dns2.com .
```

```
Member: dns3.com
```

```
Subscriber Secure Bypass enabled for the
```

```
member dns3.com .
```

```
Infoblox> set subscriber_secure_data bypass off site mobile1
```

```
Site: mobile1
```

```
Member: dns1.com
```

```
Subscriber Secure Bypass enabled for the member  
dns1.com.
```

```
Member: dns2.com
```

```
Subscriber Secure Bypass disabled for the member  
dns2.com .
```

```
Member: dns3.com
```

```
Subscriber Secure Bypass disabled for the member  
dns3.com .
```

```
Infoblox > set subscriber_secure_data bypass off
```

```
Member: dns1.com
```

```
Subscriber Secure Bypass disabled on the member dns1.com .
```

## set subscriber\_secure\_data delete

If you have configured Infoblox Subscriber Services, you can use the `set subscriber_secure_data delete` command to delete a specific subscriber record from the subscriber cache. For information about Infoblox Subscriber Services, see [Infoblox Subscriber Services](#).

### Syntax

```
set subscriber_secure_data delete <ip_addr> <prefix> <local_id>  
<ip_space_desc>
```

Argument	Description
< ip_addr >	Specify the IPv4 or IPv6 address of the subscriber.
<prefix>	Specify the prefix length.
<local_id>	Local ID present in the incoming EDNS0 packet.
<ip space desc>	The IP space discriminator. To indicate a record without an IP space discriminator, enter N/A.

## Example

```
Infoblox > set subscriber_secure_data delete 10.32.1.145 32 AABBC112233 corp1

10.32.1.145/32|LID:aabbcc112233|IPS:corp1|FLG:|ACS:Acct-Session-Id=32de327aced215ab;SUB:IMSI=301030100000026;LID:AABBC112233;NAS:NAS-PORT=168496141;PCP:Parental-Control-Policy=20001;EXP:Expire-Profile=Mon May 29 10\\:23\\:56 2017;DCP:Dynamic-Category-Policy=0;SSP:Subscriber-Secure-Policy=5F;|Thu Dec 20 10:53:41 2018

Record successfully deleted
```

## set subscriber\_secure\_data clear\_all

If you have configured Infoblox Subscriber Services, you can use the `set subscriber_secure_data clear_all` command to delete all the subscriber records from the subscriber's cache. Note that only the subscriber records with the K flag are deleted, however the subscriber records with an S flag and without any flag will not be deleted, they remain in the subscriber's cache.

## Syntax

```
Infoblox > set subscriber_secure_data clear_all
```

This command has no arguments.

## Example

```
Infoblox > set subscriber_secure_data clear_all

Deleted: 1 record
```

## set subscriber\_secure\_data garbage\_collect

You can use the `set subscriber_secure_data garbage collect` command to designate the specific member for garbage collection service. To view information about the status of the command, see [show subscriber\\_secure\\_data garbage\\_collect](#).

## Syntax

```
set subscriber_secure_data garbage_collect {on | off}
```

Argument	Description
on	Enables the member for garbage collection.
off	Disables the member for garbage collection.

## Examples

```
Infoblox > set subscriber_secure_data garbage_collect on
```

```
This member is now configured for garbage collection at 2 AM everyday. Do you want to change the scheduled time for garbage collection (y or n): y
```

```
Enter garbage collection start time [<1-12> <AM/PM>]: 3 AM
```

```
Garbage collection is scheduled at 3 AM everyday.
```

## set subscriber\_secure\_data never\_proxy

If you have configured Infoblox Subscriber Services, use the `set subscriber_secure_data never_proxy` command to set a 32-bit hexadecimal character. This character represents the list of categories to be used in the global list used to resolve DNS queries without proxying to an MSP (Multi-Services Proxy) server.

You can view the hexadecimal value of the `never_proxy` category using the `show subscriber_secure_data never_proxy` command.

## Syntax

```
set subscriber_secure_data never_proxy <category hexadecimal_character>
```

### Note

You must restart the DNS service for the hexadecimal character to be set.

For a list of the different types of parental control policies and their hexadecimal values, refer to the supplemental documentation provided by Infoblox or contact your Infoblox representative.

## Example

```
Infoblox > set subscriber_secure_data never_proxy
000fffffffffffffffffffffffffffffff01
never_proxy categories are set
!!! A RESTART of the DNS service is required before this change can take effect !!!
```



In this example, the hexadecimal character of 000ffffffffffffffffffff01 represent the category "Alcohol". Therefore, any domain related to alcohol (for example, www.beer.com, www.liquor.com) is not proxied to an MSP server. Instead, it is directly resolved by NIOS.

## set subscriber\_secure\_data persist

If you have configured Infoblox Subscriber Services, you can use the `set subscriber_secure_data persist` command to enable data persistence mode which allows static records to survive restart. For information about Infoblox Subscriber Services, see [Infoblox Subscriber Services](#).

### Syntax

```
set subscriber_secure_data persist on
```

This command turn on persistence of static subscriber records.

### Example

```
Infoblox > set subscriber_secure_data persist on
```

## set support\_access

The `set support_access` command enables and disables support access. This feature is disabled (off) by default. Enabling this feature allows Infoblox Support (Tier 3 access) to perform root level diagnostics on an appliance that is in severe distress. A special key is required to access the appliance at root level, and only Infoblox Support (Tier 3) can generate this key.

---

**Note:** Once the problem has been resolved, Infoblox recommends that you turn off this port. Any open port that is not in use can become a security risk.

---

### Syntax

```
set support_access
```

This command has no arguments.

### Example

```
Infoblox > set support_access
Enable support access (Grid level)? (y or n): y
New support access settings:
```

```
Support access enabled: Yes
```

```
Is this correct? (y or n): y
```

## set support\_timeout

The `set support_timeout` command allows you to specifying a custom timeout value for the support bundle download. For information on display of the custom timeout value, see [show support\\_timeout](#).

## Syntax

```
set support_timeout [value/off]
```

Argument	Description
value	Specify the timeout value for the support bundle download. The recommended value is between 0 to 86400 seconds, the default value is 1200.
off	The timeout value is set to its default value 1200.

## Example

```
Infoblox > set support_timeout 2600
```

## set sysName

You can use the `set sysName` command to set the FQDN (fully qualified domain name) of the appliance to allow configured SNMP management system to query the sysName value. If the appliance is an HA pair, you can use the `name2` command to set the FQDN of node 2 of the HA pair.

## Syntax

```
set sysName name1 [name2]
```

Argument	Description
name1	The FQDN of the appliance.
name2	The FQDN of node 2 of an HA pair.

## Examples

### Setting the FQDN of the appliance

Enter the following on the appliance:

```
Infoblox > set sysName eng.corp100.com
```

### Setting the FQDNs of an HA Pair

Enter the following on the appliance:

```
Infoblox > set sysName active.corp100.com passive.corp100.com
```

## set tcp\_timestamps

The **set tcp\_timestamps** command allows you to enable or disable TCP timestamps. You can view these timestamps in the traffic capture file. TCP timestamps are enabled by default. If you disable TCP timestamps, the timestamps are not displayed in the traffic capture file.

You can run the [show tcp\\_timestamps](#) command to determine whether the TCP timestamps are enabled or disabled before running `set tcp_timestamps`.

## Syntax

```
set tcp_timestamps {enable|disable}
```

## Example

```
Infoblox > set tcp_timestamps enable
```

## set temp\_license

The **set temp\_license** command generates and installs a temporary 60-day license for a fully functional NIOS appliance and IBOS (Infoblox Orchestration Server). Depending on the appliance model, the list of temporary licenses varies. To view the list of licenses installed on a NIOS appliance, use the [show license](#) command. For more information about licenses, see [Managing Licenses](#).

Infoblox supports cloud API calls to set temporary licenses for **Threat Protection (Software add-on)** and **Threat Protection Update** licenses.

To install temporary license(s) for...	Select...	Remarks
DNS and DHCP services	<b>DNSone (DNS, DHCP)</b>	Installs DNS and DHCP licenses.

To install temporary license(s) for...	Select...	Remarks
DNS and DHCP services with Grid	<b>DNSone with Grid (DNS, DHCP, Grid)</b>	Installs DNS, DHCP, and Grid licenses.
DHCP services with Grid	<b>Network Services for Voice (DHCP, Grid)</b>	Installs DHCP and Grid licenses.
NIOS services	<b>Add NIOS License</b>	Installs a NIOS license for TrinziC 2016 hardware appliances.
DNS services	<b>Add DNS Server license</b>	Installs a DNS license.
DNS Cache Acceleration	<b>Add DNS Cache Acceleration</b>	Installs a DNS Cache Acceleration license.
DHCP services	<b>Add DHCP Server license</b>	Installs a DHCP license.
Grid	<b>Add Grid license</b>	Installs a Grid license.
Microsoft server management	<b>Add Microsoft management license</b>	Installs a Microsoft server management license.
Multi-Grid management	<b>Add Multi-Grid Management license</b>	Installs a Multi-Grid license for Multi-Grid management.
Query Redirection	<b>Add Query Redirection license</b>	Installs a query redirection license.
Software Advanced DNS Protection (ADP)	<b>Threat Protection (Software add-on)</b>	Installs a Software Advanced DNS Protection license for supported platforms.
Threat protection	<b>Threat Protection Update license</b>	Installs a Threat Protection Update license.
Response Policy Zones support	<b>Add Response Policy Zones</b>	For support of Response Policy Zones.
FireEye integrated RPZ support	<b>Add FireEye license</b>	Installs a FireEye license.
DNS Traffic Control	<b>Add DNS Traffic Control license</b>	Installs a DNS Traffic Control license.

To install temporary license(s) for...	Select...	Remarks
Cloud Network Automation	<b>Add Cloud Network Automation license</b>	Installs a Cloud Network Automation license.
Cloud Platform Appliances	<b>Add Cloud Platform license</b>	Installs a Cloud Platform license.
Security Ecosystem	<b>Add Security Ecosystem license</b>	Installs a Security Ecosystem license.
Threat Analytics	<b>Add Threat Analytics license</b>	Installs a Threat Analytics license.
Flex Grid Activation	<b>Flex Grid Activation license</b>	Installs the following licenses: Grid (enterprise), Unbound, DNS Cache Acceleration (DCA), DNS, DHCP, DNS Traffic Control (DTC), Response Policy Zone, Dual Engine DNS (only for recursive DNS), Software Threat Protection (sw_tp), Threat Protection Update (tp_sub), DNSFW (rpz), NXDOMAIN Redirect (qrd), FireEye, Threat Analytics, Microsoft Management, Security Ecosystem, and Cloud Network Automation.
FLEX Grid Activation for Managed Services	<b>FLEX Grid Activation for Managed Services license</b>	Installs FLEX Grid Activation for Managed Services license. It includes the same set of licenses that are bundled with the Flex Grid Activation license.
Discovery service	<b>Add Discovery license</b>	Installs a network discovery license.
Reporting service	<b>Add Reporting license</b>	Installs a license on the reporting server.
Reporting Subscription service	<b>Reporting Subscription</b>	Installs a Reporting Subscription license, which is a Grid-wide license.



#### Note

To configure an appliance as an IB-FLEX, use the `set hardware-type` command. For more information, see [set hardware-type](#).

## Syntax

```
set temp_license
```

This command has no arguments.

## Example

### Installing a Temporary License on the NIOS Appliance

```
Infoblox > set temp_license
```

1. DNSone (DNS, DHCP)
2. DNSone with Grid (DNS, DHCP, Grid)
3. Network Services for Voice (DHCP, Grid)
4. Add NIOS license
5. Add DNS Server license
6. Add DHCP Server license
7. Add Grid license
8. Add Microsoft management license
9. Add Multi-Grid Management license
10. Add Query Redirection license
11. Add Threat Protection (Software add-on) license
12. Add Threat Protection Update license
13. Add Response Policy Zones license
14. Add FireEye license
15. Add DNS Traffic Control license
16. Add Cloud Network Automation license
17. Add Security Ecosystem license
18. Add Threat Analytics license
19. Add Flex Grid Activation license
20. Add Flex Grid Activation for Managed Services license

```
Select license (1-20) or q to quit: 4
```

1. IB-V805
2. CP-V805
3. IB-V815
4. IB-V825
5. IB-V1405
6. CP-V1405
7. IB-V1415
8. IB-V1425
9. IB-V2205
10. CP-V2205

11. IB-V2215
12. IB-V2225
13. IB-V4005
14. IB-V4015
15. IB-V4025
16. IB-V5005

Enter a number corresponding to a NIOS model (1 - 16) or q to quit:

## Installing a Temporary License on the Discovery (ND) Appliance

Infoblox > **set temp\_license**

1. DNSone (DNS, DHCP)
2. DNSone with Grid (DNS, DHCP, Grid)
3. Network Services for Voice (DHCP, Grid)
4. Add NIOS license
5. Add DNS Server license
6. Add DHCP Server license
7. Add Grid license
8. Add Microsoft management license
9. Add Multi-Grid Management license
10. Add Query Redirection license
11. Add Response Policy Zones license
12. Add FireEye license
13. Add DNS Traffic Control license
14. Add cloud Network Automation license
15. Add Security Ecosystem license
16. Add Threat Analytics license
17. Add Flex Grid Activation license
18. Add Flex Grid Activation for Managed Services license

Select license (1-18) or q to quit: 4

1. ND-V805
2. ND-V1405
3. ND-V2205
4. ND-V4005

Enter a number corresponding to a NIOS model (1-4) or q to quit:

```
Infoblox > set temp_license
```

1. Add Grid license
2. Add Discovery license

```
Select license (1-2) or q to quit:
```

## Installing a Temporary License on the Trinzic IB-V805 Reporting Appliance

```
Infoblox > set temp_license
```

1. Add NIOS license
2. Add Grid license
3. Add Reporting license

```
Select license (1-3) or q to quit:
```

## set term

Use the **set term** command to set the correct terminal type for future commands in the current session.

## Syntax

```
set term
```

This command has no arguments.

## Example

```
Infoblox > set term vt100
```

## set test\_promote\_master

Run the **set test\_promote\_master** command before promoting a Grid Master Candidate to Grid Master. The command checks whether the Grid Master Candidate is connected to the rest of the Grid members. It runs a test promotion by sending specifically crafted test packets from the Grid Master Candidate and checking whether the Grid members are able to receive these packets. To view the status of the promotion, run the [show test\\_promote\\_master](#) command.

You can run a test promotion by also using the **GM Test** option on the Grid Manager. For more information, see [Managing a Grid](#).

## Syntax

```
set test_promote_master <ip_address>
```



Argument	Description
ip_address	IP address of the Grid Master Candidate that you want to promote to Grid Master

## Examples

To test the Grid Master Candidate promotion:

```
Infoblox > set test_promote_master 10.33.151.4
```

```
Enter Timeout (>10 secs) for test messages(default 120 sec): 120
```

```
Upgrade/Downgrade is prohibited during the test. Do you want to proceed? (y or n): y
```

Here, test packets are sent every 10 seconds for 120 seconds from the Grid Master Candidate to the Grid members.

To stop a test promotion that is in progress:

```
Infoblox > set test_promote_master stop
```

```
This action will stop the promote master test which is still in progress.
```

```
Do you want to proceed? (y or n): y
```

## set thresholdtrap

Use the **set thresholdtrap** command to enable the SNMP trap for CPU usage and to configure the trigger and reset values of the trap. The CPU usage trap is disabled by default. When you use this command to change the trigger and reset values, you enable the trap and the appliance sends and resets traps based on the configured values. When CPU usage of an appliance exceeds the trigger threshold for 15 seconds, the appliance sends a "CPU usage above threshold value" trap. After the appliance sends the "CPU usage above threshold value" trap, it sends a "CPU usage OK" trap when the CPU usage dips below the reset threshold.

Use the [show thresholdtrap](#) command to view the current settings of the CPU usage trap. Note that the CPU usage trap is disabled by default, and the trigger value is set at 100 and reset value at 0. For information about Infoblox SNMP traps, refer to the *Infoblox NIOS Administrator Guide*.

## Syntax

```
set thresholdtrap {type} trigger {value} reset {value}
```

Argument	Description
type	The type of threshold trap. Enter <b>CpuUsage</b> to enable the CPU usage trap and set the trigger and reset values. Valid values are: NetworkCapacity, DBObjects, Disk, Memory, Rootfs, CpuUsage, Reporting, ReportingVolume, FDUsage
value (for trigger)	The trigger value of the SNMP trap. When CPU usage exceeds this value, the appliance sends a "CPU usage above threshold value" trap.

Argument	Description
value (for reset)	The reset value of the SNMP trap. When CPU usage dips below this value, the appliance sends a "CPU usage OK" trap. Ensure that the reset value is smaller than the trigger value.

## Example

### Enabling the CPU usage trap and set the trigger and reset values

```
Infoblox > set thresholdtrap CpuUsage trigger 80 reset 71
Infoblox >
```

When CPU usage exceeds 80% of capacity for 15 seconds, the appliance sends the "CPU usage above threshold value" trap. The appliance sends the "CPU usage OK" trap when CPU usage dips below the reset value of 71. Following is a sample SNMP output of this example:

```
2011-09-30 04:32:20 ib-10-35-107-9.infoblox.com [UDP: [10.35.107.9]:44183->[10.35.107.9]]:
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (50991) 0:08:29.91
SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.7779.3.1.1.1.1.3.0
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.1.0 = STRING: "10.35.107.9"
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.3.0 = STRING: "cpu_usage"
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.6.0 = INTEGER: 100
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.7.0 = INTEGER: 80
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.8.0 = INTEGER: 71
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.11.0 = STRING: "CPU usage above threshold
value."
2011-09-30 04:33:12 ib-10-35-107-9.infoblox.com [UDP: [10.35.107.9]:44183->[10.35.107.9]]:
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (56202) 0:09:22.02
SNMPv2-MIB::snmpTrapOID.0 = OID: SNMPv2-SMI::enterprises.7779.3.1.1.1.1.3.0
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.1.0 = STRING: "10.35.107.9"
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.3.0 = STRING: "cpu_usage"
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.6.0 = INTEGER: 5
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.7.0 = INTEGER: 80
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.8.0 = INTEGER: 71
SNMPv2-SMI::enterprises.7779.3.1.1.1.2.11.0 = STRING: "CPU usage OK."
```

## set token

The `set token` command is used to set Token, Grid Masters IP address and Certificate to the token file.

### Syntax

```
set token [on|off|join]
```

Argument	Description
on	Sets Token, Grid Masters IP address and Certificate to the token file
off	Removes Token from the file and all corresponding data
join	To join the grid using the existing token file

### Example

```
Infoblox > set token on
```

## set traffic\_capture

The `set traffic_capture` command allows you to capture the traffic for one or all of the ports on a NIOS appliance and save the traffic capture in a file. To capture traffic, the NIOS appliance must have a minimum of 500 MB of free disk space; otherwise, the traffic capture might fail.

The NIOS appliance saves all traffic it captures in a .cap file and compresses it into a .tar.gz file. The size of the .cap file is limited to 4 GB for Infoblox-4030-10GE and the size is limited to 1 GB for all other NIOS appliances. In Grid Manager, you can download the traffic capture file after the traffic capture stops by navigating to the **Grid** tab -> **Grid Manager** tab -> **Members** tab -> and click **Traffic Capture** from the Toolbar. To view information about the traffic capture running on the member, see [show traffic\\_capture\\_status](#).

You can also transfer the traffic capture file to remote hosts either using FTP or SCP. You cannot transfer the traffic capture files when the traffic capture is in progress. Note that this operation may take a long time to complete.

### Syntax

```
set traffic_capture on [port <all|lan1|lan2|mgmt|ha>] [vlan <id>] [duration  
<seconds>]  
[filter 'valid-tcpdump-filter-expression'] [with-rolling]  
set traffic_capture off  
set traffic_capture transfer [ftp|scp] <server-ip> <user-name> <user-  
password>  
[dest <file_name>]
```

Argument	Description
<code>on</code>	Starts the traffic capture.
<code>off</code>	Stops the traffic capture after you have started it.
<code>port</code> <code>  all</code> <code>  lan1</code> <code>  lan2</code> <code>  mgmt</code> <code>  ha</code>	Specifies the port for which you want to capture the traffic: <ul style="list-style-type: none"> <li>• Captures traffic on all ports.</li> <li>• Captures traffic on the LAN1 port.</li> <li>• Captures traffic on the LAN2 port.</li> <li>• Captures traffic on the MGMT port.</li> <li>• Captures traffic on the HA port.</li> </ul> The default is LAN1. Note that if you have enabled the LAN2 failoverfeature, the LAN1 and LAN2 ports generate the same output. (For information about the LAN2 failover feature, see the <i>About Port Redundancy</i> section in the <i>Infoblox NIOS Administrator Guide</i> .)
<code>vlan &lt;id&gt;</code>	Captures traffic for the vLAN interface on LAN1 or LAN2
<code>duration &lt;seconds&gt;</code>	Specifies the duration in seconds for which you want the traffic capture to run. The default is 1800 seconds.
<code>filter</code> <code>'valid-tcpdump-</code> <code>filter-expression'</code>	Allows you to set the "tcpdump" filter for traffic capture.
<code>with-rolling</code>	Enables rolling of traffic capture file. When the traffic capture file reaches the maximum size limit, the appliance automatically saves the file into a new file and continues capturing the traffic. The appliance can save up to 4 traffic capture files.
<code>transfer</code>	Allows you to transfer the traffic capture file to an FTP server or a SCP server.
<code>ftp</code>	Transfers the traffic capture file to an FTP server.
<code>scp</code>	Transfers the traffic capture file to an SCP server.

## Example

### Start the traffic capture

```
Infoblox > set traffic_capture on
Traffic capture started successfully.
```

## Stop the traffic capture

```
Infoblox > set traffic_capture off  
Traffic capture stopped successfully.
```

## Start the traffic capture on a specific port

```
Infoblox > set traffic_capture on port lan1  
Traffic capture started successfully.
```

## Specify the traffic capture filter

```
Infoblox > set traffic_capture on port lan1 filter 'host 10.32.2.34'  
Traffic capture started successfully.  
Infoblox > set traffic_capture on port lan1 filter 'net 192.168.0.0/24'  
Traffic capture started successfully.  
For information about valid tcpdump filter expressions, refer to https://  
wiki.wireshark.org/CaptureFilters.
```

## Transfer traffic capture file to an FTP server

```
Infoblox > set traffic_capture transfer ftp 10.120.20.239 frtest Infoblox123  
dest  
/home/rpadasalagi/ftp_back.tar.gz  
WARNING: This operation may take a long time to complete  
Do you want to proceed? (y or n):y
```

## Transfer traffic capture file to an SCP server

```
set traffic_capture transfer scp 10.120.20.239 frtest - dest /home/test/  
scp_back.tar.gz  
Enter password:  
WARNING: This operation may take a long time to complete  
Do you want to proceed? (y or n):y
```

---

**WARNING:** When you use the `set traffic_capture transfer` command, the user password is logged in the history. In order to avoid exposure of the user password, you can enter dash (-) instead of the actual

password in the `<user-password>` field. You can enter the password when the appliance prompts for the password.

## set txn\_trace

The `set txn_trace` command enables/disables tracing of DB transaction.

### Syntax

```
set txn_trace [on|off]
```

Argument	Description
on	Enable tracing for DB transactions
off	Disable tracing for DB transactions

### Example

```
Infoblox > set txn_trace
```

## set update\_rabbitmq\_password

The `set update_rabbitmq_password` command updates the RabbitMQ password from the unsecure to the secure mode. When you perform a staged upgrade of NIOS from any source version to version 8.6.2 or later, Infoblox recommends that you run this command to make the password secure.

#### Note

Running the `set update_rabbitmq_password` command restarts the Grid Manager. Therefore, Infoblox recommends that you run the command within the maintenance window.

### Syntax

```
set update_rabbitmq_password
```

This command has no arguments.

### Example

```
Infoblox > set update_rabbitmq_password
```

```
Switching RabbitMQ password to secure mode requires a product restart.
```

```
Are you sure you want to proceed? (y or n): y
```

```
RabbitMQ password updated.
```

## set upgrade\_dist rsync\_batch disable

The set upgrade\_dist rsync\_batch disable command disables the resync-batch mode for all members on the Grid.

### Syntax

```
set upgrade_dist rsync_batch disable
```

### Examples

```
Infoblox > set upgrade_dist rsync_batch disable
```

## set wins\_forwarding

Use the `set wins_fowarding` command to configure the appliance to forward WINS packets to Microsoft DNS and DHCP servers. You can enable or disable this feature for the entire Grid or override Grid-level settings for specific Grid members. To configure Grid level setting, you must execute this command on the Grid Master.

To view information about the current packet forwarding configuration, see [show wins\\_forwarding](#).

### Syntax

```
set wins_forwarding grid|member
```

Argument	Description
grid	Configures all members in the Grid to enable or disable the forwarding of WINS packets to Microsoft DNS and DHCP servers. Grid-level setting can be set only on the Grid Master.
member	Overrides the Grid settings to enable or disable packet forwarding for a specific Grid member.

### Examples

#### Enable packet forwarding for the Grid

```
Infoblox > set wins_forwarding grid
```

```
This command will change default settings for WINS packets forwarding (will
```

```
affect members inheriting grid settings). Continue? (y or n): y
Enable WINS packets forwarding? (y or n): y Enter default WINS server IP:
1.0.0.123 Select output interface from the list:
a. MGMT
b. LAN
c. LAN2
d. VIP

WINS packets forwarding is enabled. WINS packets will be forwarded to
"1.0.0.123".
```

### Override Grid-level packet forwarding configuration for a specific Grid member

```
Infoblox > set wins_forwarding member
WINS forwarding on grid level is enabled. WINS server IP is 1.0.0.123.
Override grid level settings? (y or n): y


Enable WINS packets forwarding on this member? (y or n): y Enter WINS
server IP: 1.0.0.321
Select output interface from the list:
1. MGMT
2. LAN
3. LAN2
4. VIP

WINS packets forwarding is enabled. WINS packets will be forwarded to
"1.0.0.321".
```

### Disable packet forwarding for the Grid

```
Infoblox > set wins_forwarding grid
This command will change default settings for WINS packets forwarding (will
affect members inheriting grid settings). Continue? (y or n): y


Enable WINS packets forwarding? (y or n): n


WINS packets forwarding was disabled on grid level.
```



## Overriding Grid-level configuration and disable packet forwarding for a specific Grid member

```
Infoblox > set wins_forwarding member
WINS forwarding on grid level is enabled. WINS server IP is 1.0.0.123.

Override grid level settings? (y or n): y
Enable WINS packets forwarding on this member? (y or n): n
WINS packets forwarding is disabled for this member.
```

## show admin\_group\_acl

The `show admin_group_acl` command displays admin groups that have ACL settings.

### Syntax

```
show admin_group_acl
```

This command has no arguments.

### Example

The following example shows the response when none of the admin groups have enabled ACL settings.

```
Infoblox > show admin_group_acl
None of Admin Groups have enabled ACL settings.
```

## show allow\_query\_domain

The `show allow_query_domain` command displays the list of all domain names in the DNS view specified or its default DNS view. For information about adding, updating, or deleting the allow query domain ACLs, see [set allow\\_query\\_domain](#). To view the list of all DNS views that have the allow query domain ACLs, see [show allow\\_query\\_domain\\_views](#).

### Syntax

```
show allow_query_domain [view-name]
show allow_query_domain <view-name> [domain-name]
```

Arguments	Description
<code>view-name</code>	Displays the named ACLs configured under the view name.
<code>domain-name</code>	Displays the domains in which the named ACLs are configured.

## Example

```
Infoblox > show allow_query_domain
```

```
DNS View: default
```

```
Domain Name : foo.com
```

```
Named ACL : Named_ACL
```

```
Domain Name : abc.com
```

```
Named ACL : Named_ACL2
```

## show allow\_query\_domain\_views

The `show allow_query_domain_views` command displays the list of all DNS views that have the allow query domain ACLs configured. For information about adding, updating, or deleting the allow query domain ACLs, see [set allow\\_query\\_domain](#). To view the list of all domains that have allow query domain ACLs, see [show allow\\_query\\_domain](#).

## Syntax

```
show allow_query_domain_views
```

This command has no arguments.

## Example

```
Infoblox > show allow_query_domain_views
```

```
DNS views that have domain ACL associations:
```

```
1. default
```

```
2. dns_view
```

## show analytics\_parameter

The `show analytics_parameter` command displays configuration of Analytics service.

## Syntax

```
show analytics_parameter [grid|member]
```

Argument	Description
grid	Displays configuration of Grid Analytics service
member	Displays configuration of member Analytics service

## Example

```
Infoblox > show analytics_parameter
```

## show adp

The **show adp** command displays ADP (Advanced Threat Protection) details on the supported platform. You can use this command only if **Threat Protection** (hardware based) or **Threat Protection (Software add-on)** licenses are installed on the platform.

## Syntax

```
show adp
```

This command has no arguments.

## Example

```
Infoblox > show adp
Threat Protection:           Enabled
Threat Protection monitor mode: Disabled
Threat Protection event stats: CRITICAL=0 MAJOR=0 WARNING=0
INFORMATIONAL=2
Log level:                   6(Info)
```

## show arp

The **show arp** command displays ARP (Address Resolution Protocol) data to view mappings. This allows you to see if the current state matches the mappings. If the mappings are out of date, use the [reset arp](#) command. This command is also useful for troubleshooting network connectivity issues.

## Syntax

```
show arp
```

This command has no arguments.

## Example

The following example shows the IP address (10.1.1.1), MAC address ( 00:04:96:1D:19:80 ), and type of connection (ethernet).

```
Infoblox > show arp
? (10.1.1.1) at 00:04:96:1D:19:80 [ether] or LAN
```

## show action\_to\_activate\_hotfix

The `show action_to_activate_hotfix` command displays the best action recommended to be performed for activating the latest installed hotfix changes. Note that this can be used only on Grid Master.

## Syntax

```
show action_to_activate_hotfix <host_name> [<nios_version>] [resync]
```

Argument	Description
host_name	Enter the hostname. Note this is mandatory.
nios_version	Enter the NIOS version. Note this is optional.
resync	This option is used to re-synchronize to the hotfix manifest file from the Grid member. Note this is optional.

## Example

```
Infoblox > show action_to_activate_hotfix infoblox.localdomain
```

```
Hotfix generic name      : CHF-8.6.3.1-J90002-apply-1681369836
Hotfix time              : 15-06-23 07:10:36 UTC
Suggested best action to activate : Node restart required
Member status           : ONLINE
```

Note: This action is to be performed after applying the hotfix, **if** already done please ignore.

## show auto\_provision

The `show auto_provision` command shows the state of auto-provisioning for an appliance. It displays whether auto-provisioning is enabled or disabled for an appliance.

### Syntax

```
show auto_provision
```

This command has no arguments.

### Example

```
Infoblox > show auto_provision
```

```
Auto Provision is enabled
```

```
Infoblox > show auto_provision
```

```
Auto Provision is disabled
```

## show bfd

The `show bfd` command displays the detailed BFD information, session details, and configuration information. For information about how to set the logging level, see [set bfd](#).

### Syntax

```
show bfd [ summary | details | config ]
```

Argument	Description
<code>summary</code>	Displays the BFD neighbor information.
<code>details</code>	Displays the BFD neighbor detailed information.
<code>config</code>	Displays the running BFD configuration file.

## Example

```
Infoblox > show bfd details
OutAddr           NeighAddr           LD/RD           Holddown(mult)
State           Int
10.34.54.68       10.34.54.16         2/4
300(3)           Up           bond0
Local Diag: 0, Demand mode: 0, Poll bit: 0
MinTxInt: 100000, MinRxInt: 100000, Multiplier: 3
Received MinRxInt: 100000, Received Multiplier: 3
Holddown (hits): 300(0), Hello (hits): 100(1638) Authentication:None
Last Sequence Number: Rx: 0, Tx: 1566182577
Rx Count: 1638
Tx Count: 1686
Last packet: Version: 1           - Diagnostic: 0
State bit: Up           - Demand bit: 0
Poll bit: 0           - Final bit: 0
Multiplier: 3           - Length: 24
My Discr: 2           - Your Discr: 4
Min tx interval: 100000- Min rx interval: 100000 Min Echo interval: 100000
```

## show bgp

The `show bgp` command displays information about the BGP configuration on the appliance, reachability information about neighbors, and BGP routes to destinations. You can specify the command with or without an argument. A command without an argument defaults to `show bgp route`.

For information about how to write statistical information to syslog, see [set bgp log](#).

## Syntax

```
show bgp {route | neighbor | summary | config}
```

Argument	Description
route	Displays the BGP routing table.
summary	Displays the BGP protocol summary.
neighbor	Displays information about all known BGP neighbors.

Argument	Description
config	Displays the running BGP configuration file.

## Example

The following examples are for illustration only. The actual output varies based on the Quagga version.

Infoblox > **show bgp route**

BGP table version is 0, local router ID is 50.0.1.2

Status codes: s suppressed, d damped, h history, \* valid, > best, i - internal, r RIB-failure, S Stale, R Removed

Origin codes: i - IGP, e - EGP, ? - incomplete

	Network	Next Hop	Metric
LocPrf	Weight	Path	
* i	10.0.1.0/24	10.0.1.1	
0	100	0	?
*>	10.0.2.99	11	
32768	?		
* i	10.0.2.0/24	10.0.1.99	
2	100	0	?
*>	0.0.0.0		
1	32768		?
* i	50.0.1.1/32	10.0.1.1	
0	100	0	i
*>	10.0.2.99	1011	
32768	?		
*	50.0.1.2/32	0.0.0.0	
1	32768		?
*>	0.0.0.0	0	
32768	i		
...			

Infoblox > **show bgp summary**

BGP router identifier 50.0.1.2, local AS number 65001 RIB entries 25, using 1600 bytes of memory

Peers 2, using 5024 bytes of memory

Neighbor	V	AS	MsgRcvd	MsgSent	TblVer	InQ	OutQ	Up/Down	State/PfxRcd
10.0.1.1	4	65001	89	85	0	0	0	01:11:27	10
100.0.2.4	4	65004	57	53	0	0	0	00:46:37	5

```
Total number of neighbors 2
```

## show bloxtools

The **show bloxtools** command displays the operational status of the bloxTools service and the usage percentage for the CPU, memory and disk resources.

### Syntax

```
show bloxtools
```

This command has no arguments.

### Example

Following is an example of the output displayed when you execute the command and bloxTools is enabled and its services are running:

```
Infoblox > show bloxtools
bloxTools status: enabled(GREEN)
CPU: 1%, Memory: 21%, Hard Disk: 0%
```

## show capacity

The **show capacity** command displays database capacity limits for your NIOS appliance. This allows you to see the object counts and types on a member. This command is useful to see the amount of data is assigned to a member and how that relates to the member's specified capacity.

### Syntax

```
show capacity
```

This command has no arguments.

### Example

```
Infoblox > show capacity
Hardware Type = IB-2000
Database Capacity = 1200000 "objects"
Objects Present = 112466 (9 percent used)
```



Count	Area	Type
16638	dns	bind_a
15000	dns	bind_cname
1000	dns	bind_mx
19392	dns	bind_ns
15501	dns	bind_ptr
836	dns	bind_soa
500	dns	bulk_host
5000	dns	dhcp_host
385	dns	dhcp_member
322	dns	dhcp_range
1538	dns	fixed_address
5000	dns	host
5000	dns	host_address
5000	dns	host_alias
265	dns	network
263	dns	shared_network_item
500	dns	srg_zone_linking
840	dns	zone
18018	dns	zone_cluster_secondary_server
537	dns	zone_ext_secondary_server
208	Grid	product_license

Note: Counts per object type not displayed unless at least 100 of that type exist.

## show cc\_mode

The `show cc_mode` command displays the Common Criteria settings.

### Syntax

```
show cc_mode
```

This command has no arguments.

## Example

```
Infoblox > show cc_mode
Common Criteria Mode Setting:
Common Criteria Mode Enabled (grid-level): Yes
```

## show certificate\_auth\_admins

The `show certificate_auth_admins` command displays whether the certificate authentication service is enabled for admins.

## Syntax

```
show certificate_auth_admins
```

This command has no arguments.

## Example

```
Infoblox > show certificate_auth_admins
Certificate authentication is enabled for next admins:
  admin
  ...
  [username]
```

## show certificate\_auth\_services

The `show certificate_auth_services` command displays the list of certificate authentication services that are used as effective authorization policies.

## Syntax

```
show certificate_auth_services
```

This command has no arguments.

## Example

```
Infoblox > show certificate_auth_services
Effective Certificate Authentication Services:
```

```
DoD CaC
```

```
[service name]
```

## show check\_auth\_ns

The `show check_auth_ns` command displays check authoritative NS RRset setting.

### Syntax

```
show check_auth_ns
```

This command has no arguments.

### Example

```
Infoblox > show check_auth_ns  
Check authoritative NS RRset is disabled
```

## show clusterd\_info

The `show clusterd_info` command displays clusterd run-time information.

### Syntax

```
show clusterd_info
```

This command has no arguments.

### Example

```
Infoblox > show check_auth_ns  
  
Waiting 5 sec while clusterd is creating dump file...  
g_am_master_vnode: true (configured as a master node)  
g_clusterd_max_nodes: 350  
g_active_GM_openvpn_serv_cnt (actual count of OpenVPN servers to run): 1  
g_first_join_attempt: false  
g_using_conn_config_file (using conn info sent by grid master): false  
g_delay_master_run (postpone starting master): false  
g_dirty_shutdown: false  
g_running_one (are the ONE services running?): true  
g_is_real_unit (true for real HW): true
```

```
g_power_down_if (power-down unused interfaces): true
g_is_vnios (true for vNIOS): false
g_udp_vrrp (true for Platforms that use UDP based VRRP): false
g_am_active_master (am I the current grid master?): true
g_am_master_vnode (am I part of grid master vnode?): true
g_was_master: true
g_ha_enabled (is HA enabled?): false
g_active (am I the active node in an HA pair?): true
g_vpn_server_setup (is the VPN server setting up): false
g_directing_upgrade (are we directing the upgrade of the grid?): false
g_reverted (a flag indicating that the member just reverted): false
g_need_grid_upgrade_state_update: false
g_upgrade_lite (lite upgrade in progress): false
g_partner_upgrade_mode: UPGRADE_NONE
g_db_locked_for_snapshot: false
g_rollback_failed (meaningful only if g_db_locked_for_snapshot is set): false
g_removing (am I being removed?): false
g_start_proxycd (should proxycd be running?): false
g_partial_replication_disabled (last known state): false
g_am_logging_member: false
g_failover_syncing: false
g_tftp_limit (grid-wide tftp storage limit, in MB): 500
g_subgrid_cnt (number of subgrids): 0
g_using_external_time_servers (is using external NTP time servers): false
```

## show config

The **show config** command displays the DNS, DHCP, DHCPv6, or DTC configuration files, named.conf, dhcp.conf, dhcpv6.conf, dtc.conf, and healthd.conf.

## Syntax

```
show config { dns | dhcp | dhcpv6 | dtc | healthd }
```

**show config** displays the contents of named.conf, dhcp.conf, dhcpv6.conf, dtc.conf, and healthd.conf files. You can page through the output 10 lines at a time.

Argument	Description
dns	Displays the named.conf file.
dhcp	Displays the dhcp.conf file.
dhcpv6	Displays the dhcpv6.conf file.
dtc	Displays the dtc.conf file at /Infoblox/var/idns_conf/dtc.conf
healthd	Displays the healthd.conf file at /Infoblox/var/idns_conf/healthd.conf

## Example

```

Infoblox > show config dns
  include "/infoblox/var/named_conf";
  options {
    zone-statistics yes;
    directory "/infoblox/var/named_conf"; version "";
    recursion no;
    listen-on {127.0.0.1; 10.0.0.0;}
    query-source address 10.0.0.0;
  }
  Enter <return> to continue with More lines or enter q<return> to go back
to the command line.

```

```

Infoblox > show config dhcp
  local-address 10.0.0.0.;
  server-identifier 10.0.0.0;
  ddns-update-style interim;
  authoritative;
  option domain-name "corp100.com";
  mini-lease-time 43200;
  max-lease-time 43200; ping-check false;
  log-facility daemon;
  Enter <return> to continue with More lines or enter q<return> to go back
to the command line.

```

```

Infoblox > show config healthd
# AUTO GENERATED FILE, DO NOT EDIT
{

```

```

'servers': [],
'monitors': [],
'checks': [],
'options': {
'source': 'VIP',
'log_idns_health': False,
'log_facility': 29,
'source_ip': '10.120.21.129'
'socket_limit': 9000
}
}

```

## show connections

The **show connections** command shows the active Internet connections for the NIOS appliance. Use this command to investigate connectivity issues or processes that may have stopped running.

### Syntax

```
show connections
```

This command has no arguments.

### Example

The following example provides information on:

- **Proto:** Active protocol, TCP or UDP
- **Recv-Q:** Packets received
- **Send-Q:** Packets sent
- **LocalAddress:** Host name and type of connection
- **ForeignAddress:** IP address of the system connected to the appliance
- **State:** State of the connection

```
Infoblox > show connections
```

```
Active Internet connections (servers and established)
```

```
Proto Recv-Q Send-Q Local Address Foreign Address State
```

```

tcp 0 0 localhost:kdm*:LISTEN
tcp 0 0 localhost:cluster-disk*:LISTEN
tcp 0 0 localhost:localdo:domain*:LISTEN
tcp 0 0 localhost:domain*:LISTEN
tcp 0 0 localhost:rndc*:LISTEN
tcp 0 0 infoblox:localdom:https*:LISTEN

```

```
tcp 0 0 localhost:https*:LISTEN
tcp 0 0 infobloxlocaldom:https10.1.1.1:arbotext-lm ESTABLISHED
Enter <return> to continue with More lines or enter q<return> to go back to
the command line.
```

## show connection\_limit

You can use the **show connection\_limit** command to display the per client IP address maximum connection limit for the following protocols: HTTP and HTTPS. Note that maximum connections here refer to the network level connections, not application level connections. For example, an HTTPS connection limit of 4 means that there can be a maximum of four TCP connections between any given client IP address and the appliance that are concerned using the HTTPS protocol.

To set the maximum connection limit, see [set connection\\_limit](#).

## Syntax

```
show connection_limit {http | https}
```

Argument	Description
http	The maximum connection limit for the HTTP protocol.
https	The maximum connection limit for the HTTPS protocol.

## Examples

### Showing the Per Client Address Maximum Connection Limit for the HTTP Protocol

```
Infoblox > show connection_limit http
Current http connection limit: 150
```

## show cpu

The **show cpu** command displays the processor and memory statistics for the NIOS appliance. This command tells you how busy the appliance is and indicates if an appliance is not performing properly.

## Syntax

```
show cpu
```

This command has no arguments.

## Example

The following example provides information on:

- **swpd**: Amount of virtual memory (swap space) used
- **free**: Amount of available memory
- **idle**: Amount of idle memory
- **buff**: Amount of memory used as buffers (I/O)
- **cache**: Amount of memory used as cache (kernel used memory)
- **swap**
  - **si**: Amount of memory paged in (per/sec) from swap or the file system
  - **so**: Amount of memory swapped out
- **io**: Disk input/output
  - **bi**: Blocks (4K each) received from a block appliance
  - **bo**: Blocks sent to a block appliance
- **system**
  - **in**: Number of hardware interrupts per second—measures how busy the CPU is
  - **cs**: Number of context switches per second—measures how busy the CPU is
- **cpu**: Measures (%) of CPU used in each of these areas—the total equals 100%
  - **us**: Percentage of CPU used running the Infoblox product and other non-kernel processes
  - **sy**: Percentage of CPU used running kernel processes
  - **id**: Percentage of CPU that is currently idle
  - **wa**: Percentage of CPU spent waiting for disk I/O

---

**Note:** If the combined io (bi and bo) and cpu wa values are high, it is a sign that the system is overloaded.

---

```
Infoblox > showcpu
```

```
-----memory----- -swap-- --io--- --system- -----cpu-----  
swpd   free   idle   buff   cache  si   so  bi  bo  in  cs  us  sy  id  wa  st  
0      984024 15432 51932 916660 0    0  0  42 61 94 3  2  95 0  0
```

## show date

The **show date** command displays the current date, time zone, and time of a NIOS appliance. Use this command if you log in to an appliance from a different time zone. This command is helpful when dealing with a Grid that comprises multiple appliances in multiple time zones.

## Syntax

```
show date
```

This command has no arguments.



## Example

```
Infoblox > show date
Tue Aug 16:24:19 EDT 2006
```

## show debug

The **show debug** command shows whether debug logging is on or off. When on, debug logging captures all traffic and processes on the NIOS appliance. Due to the verbose nature and the extent of the information captured, these log files grow at a rapid rate. For information on how to turn on or off the debug logging, see [set debug](#).

### Note

Infoblox recommends that you turn off debug logging unless Infoblox Support specifically directs you to turn on this feature. If left on, debug logging can cause performance issues.

## Syntax

```
show debug
```

This command has no arguments.

## Example

```
Infoblox > show debug
Debug logging status : enabled

Infoblox > show debug
Debug logging status : disabled
```

For information on viewing the output of the debug log file, see [show log](#).

## show debug\_analytics

The **show debug\_analytics** command shows the status of Analytics service debugging.

## Syntax

```
show debug_analytics
```

This command has no arguments.

## Example

```
Infoblox > show debug_analytics
show debug_analytics
off
```

## show delete\_tasks\_interval

The **show delete\_tasks\_interval** command shows the time interval the appliance waits until it deletes the completed and rejected tasks from the system. Once the tasks are removed, they are no longer displayed in the **Task Manager** tab of **Grid Manager**. The default interval is 14 days. For information about how to configure this time interval, see [set delete\\_tasks\\_interval](#).

## Syntax

```
show delete_tasks_interval
```

This command has no arguments.

## Example

```
Infoblox > show delete_tasks_interval
show delete_tasks_interval
Current delete tasks interval is 25 days
```

## show dhcp\_gss\_tsig

The **show dhcp\_gss\_tsig** commands provide information about an Infoblox DHCP server that is configured to send GSS-TSIG authenticated DDNS updates to an AD integrated DNS server. You can use these commands to verify your configuration and troubleshoot potential issues. You can also test whether the appliance can communicate with the Key Distribution Center (KDC) and the DNS server.

## Syntax

```
show dhcp_gss_tsig config show dhcp_gss_tsig keytab
show dhcp_gss_tsig test dns {ns-ip}
show dhcp_gss_tsig test kdc {kdc-ip} {principal}
```

Argument	Description
config	Displays configuration information and runs tests to verify the configuration.

Argument	Description
<code>keytab</code>	Displays information about the keytab file that is in use by the appliance.
<code>test dns</code>	Tests the configuration by verifying that the destination DNS server is reachable. You must enter the IP address of the destination DNS server.
<code>test kdc</code>	Tests the configuration by obtaining a test Ticket Granting Ticket (TGT) from the configured KDC. You must specify the following: <ul style="list-style-type: none"> <li>• IP address of the AD domain controller that hosts the KDC for the domain</li> <li>• The Kerberos principal you specified when you generated the keytab file, in the following format: <i>username/instance@REALM</i> <ul style="list-style-type: none"> <li>• <i>username/instance</i>: The AD user name for the NIOS appliance and the FQDN of the appliance. This entry must be the same on the AD domain controller and the Infoblox appliance.</li> <li>• <i>REALM</i>: The Kerberos realm in uppercase.</li> </ul> </li> </ul>

## Examples

### Displaying GSS-TSIG configuration information and run a test

The `show dhcp_gss_tsig config` command displays the following information:

- Whether DHCP and GSS-TSIG dynamic DNS updates are enabled on the appliance
- The system time in UTC
- Information about the KDC, principal, and domain

After it displays information, the appliance can test if it can obtain a TGT from the KDC and provides information about that transaction. Note that the TGT is for test purposes only and does not affect the data used by DHCP and DNS.

It can also test the external zones that are supposed to receive the DDNS updates as follows:

- Verify if the external zones specified in the member's network view match the member's domain
- Query the name server IP address for the address of the FQDN from the DNS principal
- Query the name server for the SOA of the zone. It displays a warning if the name server does not respond or if the result does not list the FQDN from the DNS principal as authoritative.

```
Infoblox > show dhcp_gss_tsig config
```

```
System time: Tue Oct 21 16:01:43 UTC 2008
```

```
D
```

```
HCP GSS-TSIG configuration for this member:
```

```
KDC address          10.34.123.4
```

```
KDC IP              10.34.123.4
```

```
Member principal    jdoe/anywhere@CORP100.LOCAL
```

```
Member domain      CORP100.LOCAL
```

```
GSS-TSIG           enabled
```

```
DDNS updates       enabled
```

```
DHCP service          enabled

```

Test KDC using member configuration? (y or n): y

```
Requesting TGT for jdoe/anywhere@CORP100.LOCAL from KDC 10.34.123.4...
Successfully obtained test TGT.
Credentials cache: FILE:/tmp/krb5_cache.18338
Principal: jdoe/anywhere@CORP100.LOCAL
Cache version: 4
Server: krbtgt/CORP100.LOCAL@CORP100.LOCAL Client: jdoe/
anywhere@CORP100.LOCAL
Ticket etype: arcfour-hmac-md5, kvno 2
Ticket length: 957
Auth time: Oct 21 12:00:52 2008
End time: Oct 21 13:00:52 2008
Ticket flags: initial, pre-authenticated
Addresses: addressless
Successful test. Test TGT destroyed.
```

This member is configured to update the following zones:  
corp100 on 10.34.123.4 as dns/ns-corp100.CORP100.local

Test configured zones? (y or n): y

```
Next zone is corp100 on 10.34.123.4.
Test this zone? (y or n): y
Testing external zone corp100 on NS 10.34.123.4...
DNS principal is dns/ns-corp100.corp100.local.
Derived FQDN is ns-corp100.corp100.local.
FQDN resolves to nameserver IP.
SOA for corp100 has MNAME ns-corp100.corp100.local.
Nameserver is authoritative for zone.
Zone corp100 appears valid.
```

## Displaying information about the keytab

The `show dhcp_gss_tsig keytab` command displays the current keytab.

```
Infoblox > show dhcp_gss_tsig keytab
Vno Type Principal
7 des-cbc-md5 jdoe/instance@CORP100.LOCAL
```

## Verifying the DNS server

The `show dhcp_gss_tsig test dns` command verifies the destination DNS server by performing a reverse lookup of the IP address.

```
Infoblox > show dhcp_gss_tsig test dns 10.34.123.4
Querying DNS server 10.34.123.4...
Server responded.

Probable DNS principal: dns/ns-corp100.corp100.local
```

## Obtaining a test TGT from the KDC

When you use the `show dhcp_gss_tsig test kdc` command, the appliance tries to obtain a TGT from the KDC using the specified principal. It provides information about the transaction and upon successful completion, deletes the test TGT.

```
Infoblox > show dhcp_gss_tsig test kdc 10.34.123.4 jdoe/
instance@CORP100.LOCAL
Requesting TGT for jdoe/instance@CORP100.LOCAL from KDC 10.34.123.4...
Successfully obtained test TGT.

Credentials cache: FILE:/tmp/krb5_cache.12000
Principal: jdoe/instance@CORP100.LOCAL
Cache version: 4

Server: krbtgt/CORP100.LOCAL@CORP100.LOCAL Client: jdoe/
instance@CORP100.LOCAL
Ticket etype: arcfour-hmac-md5, kvno 2
Ticket length: 957
Auth time: Oct 21 12:30:01 2008
End time: Oct 21 13:30:01 2008
Ticket flags: initial, pre-authenticated
Addresses: addressless

Successful test. Test TGT destroyed.
```

## show dhcpd\_recv\_sock\_buf\_size

The `show dhcpd_recv_sock_buf_size` displays the current DHCP receive socket buffer size. The default is 1,536 kilobytes. For information about how to set the receive socket buffer size, see [set dhcpd\\_recv\\_sock\\_buf\\_size](#).

### Syntax

```
show dhcpd_recv_sock_buf_size
```

This command has no arguments.

### Examples

#### For NIOS Appliances

```
Infoblox > show dhcpd_recv_sock_buf_size
DHCP 'dhcpd' UDP receive socket buffer size: 1500 KB
```

## show dhcpv6\_gss\_tsig

The `show dhcpv6_gss_tsig` commands provide information about an Infoblox DHCP server for IPv6 that is configured to send GSS-TSIG authenticated DDNS updates to an AD integrated DNS server. You can use these commands to verify your configuration and troubleshoot potential issues. You can also test whether the appliance can communicate with the Key Distribution Center (KDC) and the DNS server.

### Syntax

```
show dhcpv6_gss_tsig config show dhcpv6_gss_tsig keytab
show dhcpv6_gss_tsig test dns {ns-ip}
show dhcpv6_gss_tsig test kdc {kdc-ip} {principal}
```

Argument	Description
<code>config</code>	Displays configuration information and runs tests to verify the configuration.
<code>keytab</code>	Displays information about the keytab file that is in use by the appliance.
<code>dns</code>	Tests the configuration by verifying that the destination DNS server is reachable. You must enter the IP address of the destination DNS server.

Argument	Description
kdc	<p>Tests the configuration by obtaining a test Ticket Granting Ticket (TGT) from the configured KDC. You must specify the following:</p> <ul style="list-style-type: none"> <li>• IP address of the AD domain controller that hosts the KDC for the domain</li> <li>• The Kerberos principal you specified when you generated the keytab file, in the following format: <i>username/instance@REALM</i> <ul style="list-style-type: none"> <li>• <i>username/instance</i>: The AD user name for the NIOS appliance and the FQDN of the appliance. This entry must be the same on the AD domain controller and the Infoblox appliance.</li> <li>• <i>REALM</i>: The Kerberos realm in uppercase.</li> </ul> </li> </ul>

## show disk

The **show disk** command displays the disk space that is used. Use this command to verify the amount of free disk space on the NIOS and vNIOS virtual appliances. Infoblox recommends that you regularly check the available disk space. Clear off outdated logs to maintain maximum available disk space. You should not allow the disk to become completely full because this can be detrimental to the performance of the appliance.

## Syntax

```
show disk
```

This command has no arguments.

## Examples

### For NIOS Appliances

```
Infoblox > show disk
Available disk space: 207G
Disk space used: 4%
Infoblox >
```

### For vNIOS Virtual Appliances

```
Infoblox > show disk
Available disk space: 53G
Disk space used: 5%
Overall disk capacity: 120G
Infoblox >
```

## show dns

The `show dns` command displays DNS query statistics for all DNS views. It also displays the recursive cache for the specified DNS views. This command displays IDN data in punycode.

### Syntax

```
show dns {stats | cache [wait_time ntime][dns_view...]}
```

Argument	Description
<code>stats</code>	Displays DNS query statistics for all DNS views. You can also use this command to display DNS query statistics for all the DNS views of DNS cache acceleration on IB-FLEX.
<code>ntime</code>	The maximum time (from 1 to 600 seconds) to wait for the cache file to be ready.
<code>cache dns_view</code>	Specifies the DNS views for which you want to display the recursive cache.

### Example

```
Infoblox > show dns stats
success 10
referral 0
nxrrset 0
nxdomain 0
recursion 0
failure 10
```

## show dns-accel

The `show dns-accel` command displays DNS Cache Acceleration information. This command is available for:

- IB-FLEX only if the **Flex Grid Activation** license is present in the Grid.

### Syntax

```
show dns-accel
```

### Example

```
Infoblox > show dns-accel
```



```
Log level: 2(Critical)
Cache: Enabled
Minimum cached TTL: 1
Maximum cached lifetime: 86400
Cache hit count: 0
Cache miss count: 0
DNS query stats: SUCCESS=0 NXDOMAIN=0 NXRRSET=0 FAILURE=0 REFERR
AL=0
System UDP DNS query count: LAN1=0 LAN2=0 MGMT=0 HA=0
System UDP DNS response count: 0
```

```
Infoblox > show dns-accel
```

```
Log level: 5(Notice)
Cache: Enabled
Minimum cached TTL: 1
Maximum cached lifetime: 86400
Cache hit count: 0
Cache miss count: 0
DNS query stats: SUCCESS=0 NXDOMAIN=0 NXRRSET=0 FAILURE=0 REFERRAL=0
System UDP DNS query count: LAN1=0 LAN2=0 MGMT=0 HA=0
System UDP DNS response count: 0
```

## show dns-accel-cache

The **show dns-accel-cache** command enables you to view the cache for DNS Cache Acceleration. It fetches and displays new acceleration cache data. For existing files, it displays data from the previous collection, if present. This command is available for:

- IB-FLEX only if the **Flex Grid Activation** license is present in the Grid.

## Syntax

```
show dns-accel-cache
```

## Example

```
Infoblox > show dns-accel-cache
Cache is empty
```

```
Infoblox > show dns-accel-cache
```

```
Warning: This operation will temporarily disable the acceleration cache to obtain the latest cached data.
```

```
The operation should take about 15 seconds to complete.
```

```
During this period, this member's DNS query performance may be significantly reduced, and changes to the member's DNS configuration will not be permitted.
```

```
Are you sure you want to proceed with this operation (y/n): y
```

```
Cache data obtained at 2018-12-14-06:22:51
```

```
Note: In case of multiple RR's, they will be printed on their own line after the first one, aligned with commas to match correct column.
```

```
Number,Query question,Query type,Query class,AAAA
```

```
Filtered,Recursion,View,Original TTL,Remaining TTL,Usage count,Last
```

```
accessed,RR section,Type of RR,RR name,RR TTL,RR data,MX
```

```
preference,RCode,Debug<1>,Debug<2>,PCP word
```

```
1,cnn.com,A,IN,NO,0,1,3600,3588,3,3,AA,A,cnn.com,3600,151.196.31.120,,0,3,3588
```

```
,,,,,,,,,,AD,,EDNS0,,0,
```

```
0x000000000000000000000000000000004000000000000000000
```

## show dns\_gss\_tsig

The `show dns_gss_tsig` commands provide information about an Infoblox DNS server that is configured to receive GSS-TSIG authenticated DDNS updates from a DHCP server. You can use these commands for diagnostic purposes and to troubleshoot issues.

## Syntax

```
show dns_gss_tsig counters [crypto] | keytab
```

Argument	Description
counters	Displays information from the internal counters, which are non-persistent and reset to zero when services are restarted. It displays the number of TKEY (transaction key) requests the Infoblox DNS server has accepted and the number of GSS-TSIG authenticated DDNS updates received. If you specify <b>crypto</b> , the display includes the number of successful and failed attempts to establish a security by context, by encryption type.

Argument	Description
keytab	Displays information about the DNS keytab files that are in use by the appliance.

## Example

```
Infoblox > show dns_gss_tsig keytab
```

```
Vno Type Principal
```

```
4 des-cbc-md5 DNS/ns1.local.test@TEST.LOCAL
```

## show dns-over-tls-config

The `show dns-over-tls-config` command displays the DNS over TLS configuration and includes DNS over TLS servers that are listening on port 853.

## Syntax

```
show dns-over-tls-config
```

This command has no arguments.

## Example

```
Infoblox > show dns-over-tls-config
```

```
DoT listen on v4 addresses:
```

```
10.39.51.58
```

```
DoT listen on v6 addresses:
```

```
2620:10a:6000:2745::1011
```

```
DoT listen on port: 853
```

## show dns-over-tls-stats

The `show dns-over-tls-stats` command displays the statistics of DNS over TLS sessions. Statistics displayed include but are not limited to:

- Active TLS sessions
- Total number of queries received over TLS
- Total number of responses sent over TLS

## Syntax

```
show dns-over-tls-stats
```

This command has no arguments.

## Example

```
Infoblox > show dns-over-tls-stats
```

```
IP 10.39.51.58 (TLS):
```

```
rx_packets: 0
```

```
tx_packets: 0
```

```
dropped_packets: 0
```

```
max_qry_overflow_sess_drop: 0
```

```
opened_sessions: 0
```

```
closed_sessions: 0
```

```
curr_sessions: 0
```

```
IP 2620:010a:6000:2745::1011 (TLS):
```

```
rx_packets: 0
```

```
tx_packets: 0
```

```
dropped_packets: 0
```

```
max_qry_overflow_sess_drop: 0
```

```
opened_sessions: 0
```

```
closed_sessions: 0
```

```
curr_sessions: 0
```

## show dns-over-tls-status

The `show dns-over-tls-status` command displays the status of DNS over TLS service. Use this command to verify whether the feature is enabled. For steps to enable DNS over TLS, see [Configuring DNS over TLS](#).

## Syntax

```
show dns-over-tls-status
```

This command has no arguments.

## Example

```
Infoblox > show dns-over-tls-status
DoT is enabled
DoT trace is off
DoT key logging is off
Max server sockets: 128
curr server sockets: 4
curr client sockets: 0
```

## show dns\_rrl

The `show dns_rrl` command provides information about the Grid or member DNS RRL (Response Rate Limiting) settings. You can use the `set dns_rrl` command to configure the DNS RRL settings. For more information, see [set dns\\_rrl](#).

## Syntax

```
show dns_rrl [member <hostname> | view <viewname> | grid]
```



### Note

The `show dns_rrl` command accepts the member option only on the Grid Master.

Argument	Description
member <hostname>	The FQDN of the Grid member.
view <viewname>	The name of the DNS view.
grid	Show RRL settings for the Grid.

## Examples

```
Infoblox > show dns_rrl grid
Grid RRL configuration:
  responses_per_second: 200
  window: 15 (default)
  slip: 3
```

```
log_only: false (default)
```

```
Grid logging configuration:
```

```
log rate-limit: true (default)
```

## show dnstap-stats

The `show dnstap-stats` command displays the number of queries and responses sent to the destination when the dnstap log format is enabled for high performance query logging. To configure dnstap in Grid Manager, see [Capturing DNS Queries and Responses](#). To enable the dnstap log format, see [set enable\\_dnstap](#).

## Syntax

```
show dnstap-stats
```

This command has no arguments.

## Example

```
Infoblox > show dnstap-stats
```

```
Packed queries: 1
```

```
Packed responses: 1
```

```
Total records: 2
```

```
Duration connected(s): 423
```

```
Total bytes sent: 286
```

## show dnstap-status

Use the `show dnstap-status` command to view the status of the dnstap configuration. The CLI displays the status of the all the fields in the **Toggle Advanced Mode > Logging** tab. To enable or disable the dnstap status use the [set enable\\_dnstap](#) command.

## Syntax

```
show dnstap-status
```

## Example

```
Infoblox > show dnstap-status
```

```
DNSTAP is on
```

```
DNSTAP configuration overridden on a Member level.
```

```
Enable forwarding queries category of log messages using DNSTAP: Yes
```

```
Enable forwarding DNS responses category of log messages using DNSTAP: No
IP address of DNSTAP receiver to transfer captured DNS queries/responses:
7.29.121.1
DNSTAP receiver port number: 7000
```

## show docker\_bridge

The **show docker\_bridge** command displays the IP address of the current Docker bridge. You can then choose to reset the IP address by running the [set docker\\_bridge](#) command.

### Syntax

```
show docker_bridge
```

This command has no arguments.

### Example

```
Infoblox > show docker_bridge
Current Docker Bridge settings:
Bridge Gateway/CIDR: 172.17.0.1/16
```

## show doh-config

The **show doh-config** command displays the DNS over HTTPS configuration and includes DNS over HTTPS servers that are listening on port 443.

### Syntax

```
show doh-config
```

This command has no arguments.

### Example

```
Infoblox > show doh-config
DoH listen on v4 addresses:
10.39.51.58
DoH listen on v6 addresses:
2620:10a:6000:2745::1011
DoH listen on port: 443
```

## show doh-stats

The `show doh-stats` command displays the statistics of DNS over HTTPS sessions. Statistics displayed include but are not limited to:

- Active HTTPS sessions
- Total number of queries received over HTTPS
- Total number of responses sent over HTTPS

## Syntax

```
show doh-stats
```

This command has no arguments.

## Example

```
Infoblox > show doh-stats
```

```
IP 10.39.51.58
```

```
rx_queries: 0
tx_queries: 0
dropped_packets: 0
max_qry_overflow_sess_drop: 0
opened_sessions: 11
closed_sessions: 11
curr_sessions: 0
```

```
IP 2620:010a:6000:2745::1011
```

```
rx_queries: 0
tx_queries: 0
dropped_packets: 0
max_qry_overflow_sess_drop: 0
opened_sessions: 0
closed_sessions: 0
curr_sessions: 0
```

## show doh-status

The `show doh-status` command displays the status of the DNS over HTTPS service. Use this command to verify whether the feature is enabled. For steps to enable DNS over TLS, see [Configuring DNS over HTTPS](#).



## Syntax

```
show doh-status
```

This command has no arguments.

## Example

```
Infoblox > show doh-status
DoH is enabled
DoH trace is off
DoH key logging is off
Max server sockets: 128
curr server sockets: 2
curr client sockets: 0
```

## show dscp

The `show dscp` command provides information about the Grid and member DSCP values in both decimal and hexadecimal formats. You can use the `set dscp` command to configure the DSCP value. For more information, see [set dscp](#).

## Syntax

```
show dscp
```

This command has no arguments.

## Examples

For a Grid:

```
Infoblox > show dscp
Grid Level: 30 (0x1e)
Member Level:Override grid setting
                20 (0x14)
```

```
Infoblox > show dscp
Grid Level: 30 (0x1e)
Member Level:Use grid setting
```

For an independent appliance:

```
Infoblox > show dscp
DSCP:      28 (0x1c)
```

## show dtc\_geoip

The `show dtc_geoip` command provides information about the GeoIP labels that are available in the current MaxMind location database for the respective IP address. You can run this command only if you have installed the DNS Traffic Control license. For more information about DNS Traffic Control, refer to the *Infoblox NIOS Administrator Guide*.

## Syntax

```
show dtc_geoip <ip-address>
```

Argument	Description
ip-address	Valid IPv4/IPv6 address of the host.

## Examples

```
Infoblox > show dtc_geoip 54.243.36.49
```

```
Continent = North America
```

```
Country = United States
```

```
Subdivision = Virginia
```

```
Infoblox > show dtc_geoip 2607:f8b0:400a:804::1012
```

```
Continent = North America
```

```
Country = United States
```

```
Subdivision = Atlanta
```

## show enable\_match\_recursive\_only

Use the `show enable_match_recursive_only` command to view the status of the match-recursive-only option for all DNS views on a specific Grid member. For information about how to use the match-recursive-only feature, see [set enable\\_match\\_recursive\\_only](#), and also refer to the *Infoblox NIOS Administrator Guide*.

The `show enable_match_recursive_only` command reports one of three possible states:

- **True:** The DNS view is set to use the match-recursive-only setting to restrict and filter client access for the view.
- **False:** The DNS view does not use the match-recursive-only setting.
- **Inherit:** The default, where the DNS view inherits its match-recursive-only setting from the Grid.

## Syntax

```
show enable_match_recursive_only
```

This command has no arguments.

## Example

```
Infoblox > show enable_match_recursive_only
```

```
View 'default': false
```

```
View 'dnsview1': true
```

```
View 'external': inherit
```

## show extra\_dns\_name\_validations

The `show extra_dns_name_validations` command displays the status of the additional validation that takes place on host names when creating or modifying zones, subzones, and records of type A, AAAA, host record, ALIAS, CAA, MX, and NS. You can use the `set extra_dns_name_validations` command to enable or disable the additional validation on host names. For more information, see [set extra\\_dns\\_name\\_validations](#).

## Syntax

```
show extra_dns_name_validations
```

This command has no arguments.

## Example

```
Infoblox > show extra_dns_name_validations
```

```
extra_dns_name_validations option is turned off
```

## show file

The `show file` command displays specified groups and files that you can access for diagnostic purposes. You can page through the display 10 lines at a time. Use this command to view files after you enable the bloxTools Environment service.

## Syntax

```
show file {groups | group}
```

```
show file group file1 [file2 ...] [follow]
```

If you use the **show file** command without any arguments, it displays all the files that you can manage with this command. If you use the groups argument, a list of all groups is shown.

If you use the **show file** command with a *group* argument but no *file* , it displays a list of all the files in the specified group. If you use the **show file** command with *group* and *file* arguments, you can specify a real-time (live) view of the file—the same as the using tail -f arguments. You can interrupt the display by pressing **Enter**.

Argument	Description
<code>groups</code>	Displays a list of available groups.
<code>group</code>	Displays a list of files for the specified group.
<code>group file1 [file2 ...]</code>	Displays the specified (group) files.
<code>follow</code>	Displays the contents of the file live, in real-time.

## Examples

### Viewing bloxTools Files

```
Infoblox > show file bloxtools portal_access
Showing file /storage/web-portal/udata/logs/access.log
Infoblox > show file bloxtools portal_error
Showing file /storage/web-portal/udata/logs/error.log
[Thu Sep 04 11:07:59 2008] [warn] RSA server certificate CommonName (CN)
`www.infoblox.com' does NOT match server name!?
[Thu Sep 04 11:08:03 2008] [notice] Digest: generating secret for digest
authentication
...
[Thu Sep 04 11:08:03 2008] [notice] Digest: done
[Thu Sep 04 11:08:04 2008] [warn] RSA server certificate CommonName (CN)
`www.infoblox.com' does NOT match server name!?
[Thu Sep 04 11:08:04 2008] [notice] Apache/2.2.6 (Unix) mod_ssl/2.2.6
OpenSSL/0.9.8h
DAV/2 mod_perl/2.0.3 Perl/v5.8.8 configured -- resuming normal operations
[Thu Sep 04 11:11:34 2008] [warn] RSA server certificate CommonName (CN)
`www.infoblox.com' does NOT match server name!?
[Thu Sep 04 11:11:39 2008] [notice] Digest: generating secret for digest
authentication
...
```

```
[Thu Sep 04 11:11:39 2008] [notice] Digest: done
[Thu Sep 04 11:11:40 2008] [warn] RSA server certificate CommonName (CN)
`www.infoblox.com' does NOT match server name!?
[Thu Sep 04 11:11:40 2008] [notice] Apache/2.2.6 (Unix) mod_ssl/2.2.6
OpenSSL/0.9.8h
DAV/2 mod_perl/2.0.3 Perl/v5.8.8 configured -- resuming normal operations
```

```
Infoblox > show file bloxtools portal_log
```

```
Showing file /storage/web-portal/udata/logs/syslog.log
```

```
Sep 4 11:07:55 (none) kernel: Linux version 2.6.17.4 (root@buildvm2) (gcc
version
```

```
3.2.1) #1 Thu Aug 28 02:20:30 EDT 2008
```

```
Sep 4 11:07:55 (none) kernel: On node 0 totalpages: 32768
```

```
Sep 4 11:07:55 (none) kernel: DMA zone: 32768 pages, LIFO batch:7
```

```
Sep 4 11:07:55 (none) kernel: Built 1 zonelists
```

```
Sep 4 11:07:55 (none) kernel: Kernel command line: root=/dev/root
rootfstype=hostfs
```

```
rootflags=/storage/web-portal/root_fs ubdb=/storage/web-portal/swapfile
```

```
ubdc=/storage/web-portal/storagefile mem=128M eth0=tuntap,uml_htap0
```

```
con0=null con1=pts
```

```
con2=pts hostfs=/storage/web-portal
```

```
Sep 4 11:07:55 (none) kernel: PID hash table entries: 1024 (order: 10, 4096
bytes)
```

```
Sep 4 11:07:55 (none) kernel: Dentry cache hash table entries: 16384 (order:
4, 65536
bytes)
```

```
Sep 4 11:07:55 (none) kernel: Inode-cache hash table entries: 8192 (order:
3, 32768
bytes)
```

```
Sep 4 11:07:55 (none) kernel: Memory: 124372k available
```

```
Sep 4 11:07:55 (none) kernel: Calibrating delay loop... 1648.23 BogoMIPS
(lpj=8241152)
```

```
Sep 4 11:07:55 (none) syslog-ng[699]: syslog-ng starting up; version='2.0.6'
```

```
Sep 4 11:07:55 (none) kernel: Mount-cache hash table entries: 512
```

```
Enter <return> to continue with More lines or enter q<return> to proceed to
the next
```

```
file
```

## show fips\_mode

The `show fips_mode` command displays settings of the FIPS mode.

### Syntax

```
show fips_mode
```

This command has no arguments.

### Example

```
Infoblox > show fips_mode
FIPS Mode Setting:
FIPS Mode Enabled (grid-level): Yes
```

## show hardware\_status

The `show hardware_status`

command displays information about the various hardware components of a NIOS appliance. It displays the power supply status, fan speed, the CPU temperature, and status of the RAID array (for the Infoblox-2000 only).

### Syntax

```
show hardware_status
```

This command has no arguments.

### Example

The following example displays the status of an Infoblox-2000.

```
Infoblox > show hardware_status
POWER:Power OK
Fan1:5075 RPM
Fan2:4927 RPM
Fan3:4787 RPM
CPU1_TEMP: +42.0 C
CPU2_TEMP: +48.0 C
SYS_TEMP: +49 C
```

```
RAID_ARRAY: OPTIMAL
```

```
RAID_BATTERY: OK READY Yes 103 HOURS
```

The following are notes about the output:

- **POWER:** Displays the status of the power supply. The Infoblox-1552, -1552-A and -2000 have redundant power supplies. If one power supply fails, the line displays POWER FAIL. To find out which power supply failed, check the LEDs of the power supplies.
- **Fan1, Fan2, Fan3:** Displays the fan speed. The Infoblox-2000 has three fans, therefore the sample output displays the speed of each fan. If a fan is not functioning, the output displays 0 RPM.
- **CPU1 TEMP, CPU2 TEMP:** Displays the CPU temperature.
- **SYS TEMP:** Displays the operating temperature of the appliance.
- **RAID ARRAY:** Displays the status of the RAID array of an Infoblox-2000. If at least one disk is not functioning properly, this line indicates that the RAID array is degraded and lists the disks which are online. It also indicates when the RAID array is rebuilding. If there is a disk mismatch, this line indicates so and lists all the RAID disks and their disk types.
- **RAID BATTERY:** This line reports the status of the disk controller backup battery. It includes the following information:
  - **Charge status:** Displays either OK or CHARGING.
  - **Ready status:** Displays either READY YES or READY NO.
  - **Hours:** Displays the estimated number of hours remaining on the battery.



#### Note

This command is not supported on vNIOs appliances on VMware. The vNIOs appliance displays No sensors present when you enter this command.

## show hardware-type

Use the `show hardware-type` command to display the current hardware type, member host platform, and hypervisor details for your appliance.

For IB-FLEX only, you can set the hardware type on an appliance and configure it as an IB-FLEX. For more information, see [set hardware-type](#). For information about IB-FLEX, see [About IB-FLEX](#).

## Syntax

```
show hardware-type
```

This command has no arguments.

## Example

To display the hardware type of an appliance:

```
Infoblox > show hardware-type
```

```
Member hardware type: IB-V825
```

```
Member host platform = Openstack
```

```
Member hypervisor = KVM
```

### Note

The `Member hypervisor` value differs based on the underlying hypervisor. For example, in Nutanix host platforms, the `Member hypervisor` value is displayed as `HYPERV` or `KVM`. This is because Nutanix AHV supports different types of hypervisors in Windows and Linux platforms.

## show hwid

The `show hwid` command displays the hardware ID. The information provided by this command is required for acquiring a new license.

## Syntax

```
show hwid
```

This command has no arguments.

## Example

```
Infoblox > show hwid  
Hardware ID: 4dcef037e91a403fe05e10ecd241
```

## show ibtrap

The `show ibtrap` command displays whether SNMP traps and email notifications are enabled for the specified event category.

## Syntax

```
show ibtrap [category]
```

Argument	Description
<code>category</code>	Valid values are: Fan, Bloxtools, Disk, Memory, CPU, MGM, HSM, Login, PowerSupply, FTP, TFTP, HTTP, NTP, DNS, DHCP, RootFS, Database,R AID, HA, MSServer, Backup, Clear, SNMP, LCD, SSH, SerialConsole, ENAT, Network, Cluster, Controld, OSPF, IFMAP, BGP, CaptivePortal, DuplicateIP, License, System, Syslog, DiscoveryConflict, ReportingVolume, DisconnectedGrid



## Example

```
Infoblox > show ibtrap Fan
Trap Category: Fan snmp: true
email: false
```

## show interface

The `show interface` command displays network interface details. The information reveals how the NIOS appliance is connected to the network. It shows line rate, broadcast address, and whether packets are being dropped. This information allows you to check the status, find the MAC address of an appliance, and provides statistics on the quality of the network signal. This command also displays whether IPv6 is enabled. On the Infoblox-250, -550-A, -1050-A, -1550-A, -1552-A, -2000, and -2000-A appliances, the appliance displays information about the LAN2 port as well. It also displays the bonded interface information when NIC bonding is enabled in the NIOS appliance. For information about how to change your interface settings, see [set interface](#).

## Syntax

```
show interface [name | all]
```

Argument	Description
<code>name</code>	Displays information about a specific interface. Enter one of the following: lan, lan2, ha, mgmt, or loopback.
<code>all</code>	Displays information about all interfaces.

## Example

The following example illustrates how you can use the `show interface` command to view the IP address and MAC address of an appliance, and its network connection details. Note that when you manually configure the speed and duplex of an interface, the appliance communicates only the settings. When you configure the Speed/Duplex settings at 1000M/Full, auto-negotiating is required and the appliance displays `Enabled (Speed/Duplex configured by user)` instead of `Enabled` in the **Negotiation** field, as shown in the following example. Statistical information is also provided on the packets received and transmitted, as well as any errors that have occurred. Lack of packet activity can be a sign of connectivity problems, dropped packages, overruns, or collisions.

```
Infoblox > show interface
LAN:
IP Address: 10.34.33.11      MAC Address: 00:30:48:98:63:AD
Mask:255.255.255.0         Broadcast: 10.34.33.255
MTU: 1500                  Metric: 1
```

IPv6 Link: fe80::230:48ff:fe98:63ad/64  
IPv6 Status: Enabled  
Negotiation: Enabled  
Speed: 1000M Duplex: Full  
DSCP Value: 30  
Status: UP BROADCAST RUNNING MULTICAST

Statistics Information

Received

packets: 24812 bytes: 11660993 (11.1 Mb)  
errors: 0 dropped: 0  
overruns: 0 frame: 0

Transmitted

packets: 23148 bytes: 11493844 (10.9 Mb)  
errors: 0 dropped: 0  
overruns: 0 carrier: 0

Collisions: 0 Txqueuelen: 1000

LAN2:

IP Address: 10.1.1.35 MAC Address: 00:30:48:98:63:AF  
Mask: 255.255.255.0 Broadcast: 10.1.1.255  
MTU: 1500 Metric: 1  
IPv6 Link: fe80::230:48ff:fe98:63af/64  
Negotiation: Enabled  
Speed: 1000M Duplex: Full  
Status: UP BROADCAST RUNNING MULTICAST

Statistics Information

Received

packets: 11 bytes: 836 (836.0 b)  
errors: 0 dropped: 0  
overruns: 0 frame: 0

Transmitted

```
packets: 0          bytes: 0 (0.0 b)
errors:  0          dropped: 0
overruns: 0        carrier: 0
Collisions: 0      Txqueuelen: 1000
```

```
Enter <return> to continue with More lines or enter q<return> to go back
to command line
```

## show ip\_rate\_limit

The **show ip\_rate\_limit** command displays the current rate limiting rules. You configure rate limiting rules to limit access or block connections from external sources. The rules take effect immediately when you enable rate limiting. For information on rate limiting and on how to configure rate limiting rules, see [set ip\\_rate\\_limit](#).

### Syntax

```
show ip_rate_limit
```

This command has no arguments.

### Example

#### Viewing the current rate limiting rules

```
IP rate limiting is enabled.
Source          Limit          Burst
=====
10.10.1.1       0 packets/minute  0 packets
10.10.1.2       5 packets/minute  5 packets
10.10.2.1/24    5 packets/minute  10 packets
all             5000 packets/minute 5000 packets
```

## show ipv6\_bgp

The **show ipv6\_bgp** command displays the local NIOS appliance's IPv6 BGP configuration, reachability information about neighbors, and BGP routes to destinations. You can specify the command with or without an argument. A command without an argument defaults to **show bgp route**.

For information about how to write statistical information to syslog, see [set bgp log](#).

## Syntax

```
show ipv6_bgp {route | neighbor | summary | config}
```

Argument	Description
route	Displays the BGP routing table.
summary	Displays the BGP protocol summary.
neighbor	Displays information about all known BGP neighbors. If only IPv4 BGP information is available, this command option displays IPv4 information.
config	Displays the running BGP configuration file, including all IPv4 and IPv6 configuration, prefix lists and access-lists.

If no applicable information is available (for example, the current device has no IPv6 BGP configuration and hence no IPv6 BGP neighbor information), you are returned to the NIOS CLI prompt.

## Example

```
Infoblox > show ipv6_bgp summary
```

```
BGP router identifier 10.34.1.179, local AS number 1 RIB entries 3, using 288 bytes of memory
```

```
Peers 1, using 4560 bytes of memory
```

```
Neighbor    V    AS  MsgRcvd MsgSent  TblVer  InQ  OutQ  Up/Down  State/PfxRcd
10.34.1.178 4    10   192542  240631      0    0    0 02:45:16      0
```

```
Total number of neighbors 1
```

```
Infoblox > show ipv6_bgp neighbor
```

```
BGP neighbor is 2001:1938:BA8::22AA:1, remote AS 10, local AS 1, external link
```

```
BGP version 4, remote router ID 10.36.1.66
```

```
BGP state = Established, up for 02:11:21
```

```
Last read 14:34:06, hold time is 16, keepalive interval is 4 seconds
```

```
Neighbor capabilities:
```

```
4 Byte AS: advertised and received
```

```
Route refresh: advertised and received(old & new)
```

```
Address family IPv4 Unicast: advertised and received
```

Address family IPv6 Unicast: advertised and received

Message statistics: Inq depth is 0 Outq depth is 0

Inq depth is 0

Outq depth is 0

	Sent	Rcvd
Opens:	118	25
Notifications:	13	88
Updates:	219	0
Keepalives:	239773	192022
Route Refresh:	0	0
Capability:	0	0
Total:	240123	192135

Enter <return> for next page or q<return> to go back to command line.

Minimum time between advertisement runs is 30 seconds

For address family: IPv4 Unicast

NEXT\_HOP is always this router

Community attribute sent to this neighbor(both)

Inbound path policy configured

Outbound path policy configured

Incoming update prefix filter list is \*DEFAULT

Outgoing update prefix filter list is \*ANYCAST

0 accepted prefixes

For address family: IPv6 Unicast

Community attribute sent to this neighbor(both)

0 accepted prefixes

Connections established 73; dropped 72

Last reset 02:13:50, due to BGP Notification send

Local host: 10.34.1.179, Local port: 179

Foreign host: 10.34.1.178, Foreign port: 43135

Nexthop: 10.34.1.179

```
Nexthop global: 2001:db8:a22:1b0::3
Nexthop local: fe80::230:48ff:febc:9503
BGP connection: non shared network
Read thread: on Write thread: off
```

## show ipv6\_disable\_on\_dad

The **show ipv6\_disable\_on\_dad** command displays whether IPv6 is disabled, when a duplicate IPv6 address is detected, on the corresponding interface.

### Syntax

```
show ipv6_disable_on_dad
```

### Examples

```
Infoblox > show ipv6_disable_on_dad
Disable IPv6 if duplicate IPv6 address detected: off
```

```
Infoblox > show ipv6_disable_on_dad
Disable IPv6 if duplicate IPv6 address detected: on
```

## show ipv6\_neighbor

The **show ipv6\_neighbor** command displays the status, IPv6 address and link-local address (normally, the MAC address of the neighboring port) of the IPv6 neighbor for the specified NIOS appliance interface—LAN, LAN2 or MGMT.

### Syntax

```
show ipv6_neighbor <lan|lan2|mgmt>
```

If no applicable information is available (for example, the current device has no IPv6 configuration and hence no IPv6 neighbor information), you are returned to the NIOS CLI prompt.

### Example

```
Infoblox > show ipv6_neighbor lan
fe80::204:96ff:fe1d:1980 lladdr 00:04:96:1d:19:80 router STALE
```

## show ipv6\_ospf

The **show ipv6\_ospf** command displays configuration and statistical information about the OSPFv3 protocol (if any) running on the NIOS appliance. For information on changing OSPF log settings, see [set ipv6\\_ospf](#).

### Syntax

```
show ipv6_ospf {route | interface | database | neighbor | configuration}
```

Argument	Description
route	Displays the OSPF routing table, as determined by the most recent SPF calculation.
interface	Displays the state and configuration on all interfaces configured with OSPF.
database	Displays all OSPF database information.
neighbor	Displays the OSPF neighbor information.
configuration	Displays the running OSPF configuration file.

### Examples

```
Infoblox > show ipv6_ospf
```

```
OSPFv3 Routing Process (0) with Router-ID 10.34.1.179
```

```
Running 11d03:14:41
```

```
Number of AS scoped LSAs is 2
```

```
Number of areas in this router is 1
```

```
Area 0.0.0.61
```

```
Number of Area scoped LSAs is 4
```

```
Interface attached to this area: eth1
```

```
Infoblox > show ipv6_ospf interface
```

```
eth1 is up, type BROADCAST
```

```
Interface ID: 11
```

```
Internet Address:
```

```
inet : 10.34.1.179/29
```

```
inet6: 2001:db8:a22:1b0::3/64
```

```
inet6: fe80::230:48ff:febc:9503/64
```

```
Instance ID 0, Interface MTU 1500 (autodetect: 1500)
MTU mismatch detection: enabled
Area ID 0.0.0.61, Cost 1
State BDR, Transmit Delay 1 sec, Priority 1
Timer intervals configured:
  Hello 10, Dead 40, Retransmit 5
DR: 255.1.1.1 BDR: 10.34.1.179
Number of I/F scoped LSAs is 2
  0 Pending LSAs for LSUpdate in Time 00:00:00 [thread off]
  0 Pending LSAs for LSAck in Time 00:00:00 [thread off]
eth2 is down, type BROADCAST
```

## show lcd

The `show lcd` command displays whether the LCD keys are turned on or off.

## Syntax

```
show lcd
```

This command has no arguments.

## Example

```
Infoblox > show lcd
No LCD present
```

## show lcd\_info

The `show lcd_info` command displays status, CPU temperature, network settings, version, hardware ID, and licensing information. It also displays the bonded interface information when NIC bonding is enabled in the NIOS appliance. This command combines the output of the following CLI commands:

```
show status, show temperature, show network, and  
show license
```

## Syntax

```
show lcd_info
```



This command has no arguments.

## Example

```
Infoblox > show lcd_info
Grid Status: ID Grid Master
HA Status:   Not Configured

CPU_TEMP: +45.5 C

Current network settings:
IP Address:      10.35.0.20
Network Mask:    255.255.255.0
Gateway Address: 10.35.0.1
HA enabled:      false
Grid Status:     Master of Infoblox Grid

Version          : 4.2r2-0
Hardware ID      : 553a25c34f45e4a2a2349d996ae1285

License Type     : Grid
Expiration Date  : 05/31/2009
License String   : GQAAAL8oY9e0uaH3MMKfPdLXrWDTs5D4p3UerF8=

License Type     : DNS
Expiration Date  : Permanent
License String   : EQAAAL4oZM7r+K+zctvOPdLUpH3V

License Type     : Grid Maintenance
Expiration Date  : 05/31/2009
License String   : GwAAALEjbsGypr37HNSWfNLV4C/Q/5Lw6yxQ/V7Pdg==

License Type     : NIOS Maintenance
Expiration Date  : 05/31/2009
```

```
License String : GwAAALQveMGZuqbuM8iReNLV4C/Q/5Lw6yxQ/lvMJg==
```

```
License Type : DHCP
```

```
Expiration Date : Permanent
```

```
License String : EgAAAL4udMK650LibpafIZ+Y/S6A5Q==
```

 **Note**

This command is not supported on vNIOS appliances on VMware.

## show lcd\_settings

The `show lcd_settings` command displays the value of these fields: lcd\_autodim/lcd\_brightness/lcd\_hwident.

### Syntax

```
show lcd_settings
```

This command has no arguments.

### Example

```
Infoblox > show lcd_settings
```

```
LCD settings can not be configured.
```

## show license

The `show license` command displays information about the licenses installed on NIOS appliances, vNIOS appliances, and Orchestration Servers. For information, refer to the *Infoblox NIOS Administrator Guide*. You can use this command to view licenses that were transferred from one vNIOS on VMware appliance to another. For information on how to set your licenses, see [set license](#).

This command displays Grid-wide licenses when you specify gridwide or all arguments. Without any arguments, the static licenses for the member are displayed.

### Syntax

```
show license [all | csv | gridwide | revoked]
```

Argument	Description
all	Displays all product licenses in a Grid.
csv	Displays all product licenses in a Grid in csv format.

Argument	Description
gridwide	Displays Grid-wide licenses in a Grid.
revoked	Displays vNIOS licenses that were transferred from one vNIOS appliance to another.

## Examples

### Showing product licenses installed on a NIOS appliance

```
Infoblox > show license
```

```
Version      : 4.3r2-5
```

```
Hardware ID   : 6ddd3618a43027fdbb3b3ca9a29077a7
```

```
License Type  : NIOS Maintenance
```

```
Expiration Date : 05/13/2009
```

```
License String : GwAAAAsiM/VsmcoYLHostc8f6T7L7HDdR+HeL6U1WA==
```

```
License Type  : Grid
```

```
Expiration Date : 05/13/2009
```

```
License String : GQAAAAA\KONBms0BL3Ai9M8bpnPKoHLUC+HEfPM=
```

```
License Type  : DNS
```

```
Expiration Date : Permanent
```

```
License String : EQAAAAELL/oe28NFbWlz8M8arW7M
```

```
License Type  : DHCP
```

```
Expiration Date : Permanent
```

```
License String : EgAAAAEjP/ZPx44UcSQi7IJW9D2Z9g==
```

```
License Type  : Grid Maintenance
```

```
Expiration Date : 05/13/2009
```

```
License String : GwAAAA4uJfVHhdENA2Yrsc8f6T7L7HDdR+HeLfe2UQ==
```

## Showing DCA licenses installed on a NIOS appliance

```
Infoblox > show license
Version : 8.5.0-388388
Hardware ID: 4015201710000011

License Type: DNS
Expiration Date: Permanent
License String : EQAAAA4P26wTb1p3eRM8dxnr7nH8

License Type: Grid
Expiration Date: 12/27/2019
License String : GgAAAA8P3LVMLlQz0wptd1en5Ce0NEBoc5DCGtFx

License Type: NIOS (Model IB-V4025)
Expiration Date: 10/28/2019
License String : GwAAAAQIx6NCbBQmeIdtdFXq72z5IkRpeNnURtN/cA==

License Type: DNS Cache Acceleration (Tier 4)
Expiration Date: Permanent
License String : FQAAAA4CyawTb1p3eRM8dxnq8iTszxVvLg==
```

## Showing product licenses installed on a NIOS Discovery Probe appliance

```
Infoblox > show license

Version      : 6.9.0-271002
Hardware ID  : 564d5d736f92734270264e24bd7f34ea

License Type : Grid
Expiration Date : 12/21/2015
License String : GgAAALb+k/nMu+ts7UIw7sK1+7B70RJDDrQZXLr8

License Type : vNIOS (model ND-V1405)
Expiration Date : 12/21/2015
License String : HAAAAX+jvPnt6sx4hV9oMT5+LJ70gZDCfabH0Um4tA=
```

```
License Type      : Discovery
Expiration Date   : 12/21/2015
License String    : GQAAALf5lP/Rvfx351t+6Ir7+P010xNBQf6cCeU=
```

## Showing product licenses installed on an Infoblox Orchestration Server

```
Infoblox > show license
```

```
Version          : 5.x-86034
Hardware ID      : 6ddd3618a43027fdbb3b3ca9a29077a7
```

```
License Type     : IF-MAP Service
Expiration Date   : 05/04/2010
License String    : FAAAAAwTMedDlo5fICEi84MbpXPPpHCI
```

## Showing product licenses of all Grid members in csv format

```
Infoblox > show license csv
```

```
public_ip,license_type,exp_date,license_string
10.0.0.18,DNS,Permanent,EQAAAAKS4n90WFGNUSirwvyUT9/z
10.0.0.18,Grid
Maintenance,05/21/2009,EwAAAA4uJfVHhdENA2Yrsc8b6T3J7HDdR+HeLfq2Cq==
10.0.0.18,Grid,05/21/2009,GwAAAA2Z6HAtBkPFPyfg/yVRsLzI2x0kYyKaPb22g==
10.0.0.18,NIOS Maintenance,01/21/2009,GwAAAAiV/nAGGljQEDv0h/
yVRsLzI2x0kYyKb/P20Q==
10.0.0.18,DHCP,Permanent,EgAAAAEjP/ZPx44UcSQi7JW9D2Z9g==
10.0.0.22,DNS,05/21/2009,EgAAAAKU8nMlRBzcTWX63rHYFoyM0Q==
10.0.0.22,Grid
Maintenance,05/21/2009,GwAAAA4uJfVHhdENA2Yrsc8b6T3J7HDdR+HeLfq2Cq==
10.0.0.22,Grid,05/21/2009,GwAAAA2B6CftBkPFPyfg/yVRsLzI2x0kZyKaPb22g==
10.0.0.22,NIOS Maintenance,05/21/2009,GwAAAAiV/nAGGljQEDv0h/
yVRsLzI2x0kYyKb/P20Q==
10.0.0.22,DHCP,Permanent,EgAAAAEjP/ZPx96UcSQi7JW9D2Z9g==
```

## Showing vNIOS on VMware licenses that were transferred

```
Infoblox > show license revoked
```

Public IP	License Type	Exp Date	Replaced Hardware ID
10.34.196.221	Grid	Permanent	564dc31965c24cc2eb7ab2955e10e1c0

License String

PQAAADUCMoaGagzzTP0jgMU8FjNONq8dY2Ux527eLxDjCxyqsaL3woZgtPdEzhTgV+4Xk+OEIvmV  
Wk3rUf9s1Q

=

10.34.196.221 vNIOS (550) Permanent564dc31965c24cc2eb7ab2955e10e1c0  
AgAAACYCL4yHZ10rQ7vu0dLxRzAWLqtRdXQ39z+LKEW5DhjrrseLjgwf9bZEz0L1ljkWkqOEN9z  
QD4wsRU=

## Showing product licenses, including a transferred vNIOS license

Infoblox > **show license**

Version : 5.1r4-111576-ul

Hardware ID : 564d636db7a4892b1065c1d9493673a4

License Type : DHCP

Expiration Date : Permanent

License String : EgAAADQEJZOIN0/mEqkjgMU8WjBTfQ==

License Type : DNS

Expiration Date : Permanent

License String : EQAAADQCNZ/ZKwK3DuRynIhwa2MG

License Type : Grid

Expiration Date : Permanent

Revoked Hardware ID : 564dc31965c24cc2eb7ab2955e10e1c0

License String :

PQAAADUCMoaGagzzTP0jgMU8FjNONq8dY2Ux527eLxDjCxyxqsaL3woZgtPdEzhTgV+4Xk+OEIwm  
VWk3rUf9s1Q=

License Type : MS Management

Expiration Date : Permanent

License String : GwAAAD0fGY6VdB/9WvU6w4A8FjNONq8dfHB9pm+JeA==

License Type : NIOS Maintenance

Expiration Date : Permanent

License String : GgAAAD4FKZCraQvqT/ct2YhtCn4fKuJMYD1l9T7Z

License Type : vNIOS (550)

Expiration Date : Permanent

Revoked Hardware ID : 564dc31965c24cc2eb7ab2955e10e1c0

License String :

OgAAACYCL4yHZl0rQ7Vu0dLxRzAWLqtRdXQ39z+LKEW5DhjrrseLjgwf9bZEz0L1ljKwkq0EN9z  
QD4wsRU=

vNIOS: CPU cores detected: 1 - [License allows: 1]

vNIOS: CPU frequency detected: 1200MHz - [License allows: 1200MHz]

vNIOS: System memory detected: 2048MB - [License allows: 2048MB]

## Showing all product licenses in a Grid

Infoblox > **show license all**

Public IP	License Type	Kind	Exp Date	Replaced
Hardware ID				
License String				
	Security Ecosystem Grid-wide		Permanent	
HQAAALsak0zDKirMdaUsG2Yfk/j0BkhoFjhVfEtu36dJ				
10.34.12.200	Grid	Static	01/27/2017	
GQAAAN7S+0j6JImWMzxZc8VyGvYoQJyH2i60L3Y=				
10.34.12.200	DHCP	Static	Permanent	
EgAAAN/U7/30ecqDbWhZb4g8TLh7Fg==				
10.34.12.200	DNS	Static	Permanent	
EQAAAN/S//G1ZYfScSUIc8VwFesu				
10.34.12.210	DNS	Static	Permanent	
EQAAAN0m29pKq23n6iHHLriWhVRW				
10.34.12.210	DHCP	Static	Permanent	
EgAAAN0gy9YbtyC29myWMvXa3AcD6Q==				
10.34.12.210	Grid	Static	Permanent	
GgAAANwm3MMV6m0jqDiWLvXajwIevdQ7U0THmLU0				
10.34.12.220	Grid	Static	11/28/2016	
GgAAAGq4nPRvK7i2S03o7qMk9vaokCqkg6eUcc/g				
10.34.12.220	vNIOS (model ND-V1405)	Static	11/28/2016	
HAAAAHm4gf5uJ/jsRBqsoKJo9fiokz6kgeuWY57v02w=				
10.34.12.220	Discovery	Static	11/28/2016	
GQAAAGu/m/JyLa+tQVSm7+xq/LLmkuijye0TdsW=				

## Showing Grid-wide licenses in a Grid

```
Infoblox > show license gridwide
License Type      Exp Date      License String
Security Ecosystem Permanent      HQAAALsakOzDKirMdaUsG2Yfk/
j0BkhoFjhVfEtu36dJ
```

## Showing Grid-wide licenses in a Grid when an Flex Grid Activation License is installed

```
Infoblox > show license
Version          : 8.0.0-347398
Hardware ID      : 0800201605040013

License Type     : Grid
Expiration Date  : 04/20/2017
License String   : GQAAADmh7ID3wf0q0e98xJnJ79mgkh004FM2wrE=

License Type     : DHCP
Expiration Date  : Permanent
License String   : EgAAADin+5X5nL4/Z7t83dSHvpfzxA==

License Type     : DNS
Expiration Date  : Permanent
License String   : EQAAADih65mogPNue/YtwZnL58Sm

Infoblox > show license all
Public IP License Type Kind Exp Date Replaced Hardware ID License String
Flex Grid Activation Grid-wide 02/19/2017
JAAAAPwgn32cIJAtloBgYTchXdVN71rdVRT01cjljz0xvn9gygAz2g==
10.35.5.176 Grid Static 04/20/2017 GQAAADmh7ID3wf0q0e98xJnJ79mgkh004FM2wrE=
10.35.5.176 DHCP Static Permanent EgAAADin+5X5nL4/Z7t83dSHvpfzxA==
10.35.5.176 DNS Static Permanent EQAAADih65mogPNue/YtwZnL58Sm
10.35.105.10 Grid Static 02/19/2017
GgAAA0gU19juLjevCfnmAfIKoTGz4RzrxNR2mjdo
10.35.105.10 vNIOS (model IB-VM-1410) Static 02/19/2017
GgAAAPsUytLvInTyBq2jTPRGoT+z4gjpwowgnz5g
```



```
10.35.105.10 Threat Protection (Software add-on) Static 02/19/2017
FQAAAP4N/MnsInTyBq2jTPRGoT+z50i8xQ==
10.35.105.10 Threat Protection Update Static 02/19/2017 FgAAAPkK/
M7pPDn3TuCrCbpEoDn4r0roLZg=
10.35.105.10 DHCP Static 02/19/2017 FAAAA0kSwM3gb3G6S6XmAvZHp3T54xvo
10.35.105.10 DNS Static 02/20/2017 EwAAA0kU0MGtajn0SuCoAPdB7Guu5U0=
Infoblox >
```

## Showing all licenses in a Grid when Software ADP is installed

```
Infoblox > show license
Version          : 8.1.0-348290
Hardware ID      : 564d6d00229a6cd6d197ffcd1383e37b

License Type     : Grid
Expiration Date  : 03/10/2017
License String   : GgAAAN8Cp2mr0u/Es9xNAGCAeAvdy+7J5L/704mo

License Type     : vNIOs (model IB-VM-1410)
Expiration Date  : 03/10/2017
License String   : GgAAAMwCum0q3qyYvIgBTWfMeAzdyPrL4ueth4Cm

vNIOs: CPU cores detected: 4 - [License allows: 4]
vNIOs: System memory detected: 8192MB - [License allows: 8192MB]

License Type     : Threat Protection (Software add-on)
Expiration Date  : 03/10/2017
License String   : FQAAAMkbjHip3qyYvIgBTWfMeAzdzb2etQ==

License Type     : Threat Protection Update
Expiration Date  : 03/10/2017
License String   : FgAAAM4cjH+sw0Gc9cUAASn0eQqWhbiatqU=

License Type     : DNS
Expiration Date  : 03/10/2017
License String   : EwAAAN4CoHDoL+Gc8MUDAWTLNVjAz78=

License Type     : DHCP
Expiration Date  : 03/10/2017
License String   : FAAAA4EsHyLk6jR8YlNA2XNfkeXwOnJ
```

## show license\_uid

The `show license_uid` command displays the license UID of the Grid. The UID is required when requesting Grid-wide licenses. The UID that the appliance returns is the same as the **License Pool Container UID** that is used for obtaining dynamic licenses.

### Syntax

```
show license_uid
```

This command has no arguments.

### Examples

```
Infoblox > show license_uid
The grid-wide license unique ID (same as LPC_UID):
e51f90527dce4708bc1ada576286d26a
```

## show license\_pool\_container

The `show license_pool_container` command displays the license UID that is required when obtaining dynamic licenses for vNIOS virtual appliances.

### Syntax

```
show license_pool_container
```

This command has no arguments.

### Examples

```
Infoblox > show license_pool_container
The Unique ID of the License Pool Container (LPC_UID):
e51f90527dce4708bc1ada576286d26a
```

## show log\_guest\_lookups

Use the `show log_guest_lookups` CLI command to disable the logging of guest lookups. That is, this command is used to enable or disable the logging of DNS lookup requests by guest devices behind a home router.

### Syntax

```
set log_guest_lookups [on|off]
```

## show log\_txn\_id

The `show log_txn_id` command displays whether DHCP transaction ID logging is on or off. By default, DHCP transaction ID logging is enabled. Use the [set log\\_txn\\_id](#) to enable or disable logging of DHCP transaction IDs.

### Syntax

```
show log_txn_id
```

This command has no arguments.

### Example

```
Infoblox > show log_txn_id
DHCP Transaction id logging turned OFF
```

## show lom

The `show lom` command displays the LOM (Lights Out Management) settings for the IPMI interface. To configure the network settings for the IPMI interface, use the [set lom](#) command.

### Syntax

```
show lom
```

This command has no arguments.

### Example

```
Infoblox > show lom
LOM for Grid: enabled
LOM for member: enabled (inherit)

Network settings:
IP Address: 10.34.10.42
Subnet Mask: 255.255.255.0 Default Gateway IP: 10.34.10.1

Users:
```

## show max\_recursion\_depth

The `show max_recursion_depth` command displays the maximum recursion depth value.

### Syntax

```
show max_recursion_depth
```

This command has no arguments.

### Example

```
Infoblox > show max_recursion_depth  
Recursion depth limit: 7
```

## show max\_recursion\_queries

The `show max_recursion_queries` command displays the maximum recursion queries value.

### Syntax

```
show max_recursion_queries
```

This command has no arguments.

### Example

```
Infoblox > show max_recursion_queries  
  
show max_recursion_queries  
Recursion queries limit: 150
```

## show memory

The `show memory` command displays memory statistics on used and available buffers and cache. Poor performance can be an indicator that the memory is full. If your NIOS appliance is not performing as it should, use this command to verify whether or not the appliance is experiencing a memory problem. If so, Infoblox recommends that you call Infoblox Support.

### Syntax

```
show memory
```

This command has no arguments.

## Example

```
Infoblox > show memory
```

	total	used	free	buffers	cached
Mem:	1032852	309904	722948	32864	242060
Swap:	2047992	0	2047992		
Total:	3080844	309904	2770940		

## show mld\_version

The `show mld_version` command displays the version of the MLD (Multicast Listener Discovery) protocol that is running on the appliance. The appliance runs MLD version2 by default, but you can enable it to run MLD version 1 instead. For information, see [set mld\\_version\\_1](#). Note that MLDv2 is interoperable with MLDv1.

## Syntax

```
show mld_version
```

This command has no arguments.

## Example

```
Infoblox > show mld_version
```

```
Current Multicast Listener Discovery Setting:
```

```
MLD Version: 2
```

## show monitor

The `show monitor` command displays current network monitoring data, when network monitoring for DNS is turned on. This command also provides information on the average latency of authoritative and non-authoritative replies to DNS queries. Latency is the time it takes for a packet to cross a network connection, from sender to receiver.

### Note

You must turn on network monitoring for DNS to view this data. For more information, see [set monitor dns](#).

## Syntax

```
show monitor
```

This command has no arguments.

## Examples

The following example for Network Monitoring for DNS shows information on the interval times in minutes, the latency (in milliseconds), and the number of queries.

### Viewing network monitoring for DNS data

```
Infoblox > show monitor
```

```
Network Monitoring for DNS is ON
```

```
Data last updated: Tue Sep 12 19:05:51 2006
```

Authoritative	Interval (min)	Latency (usec)	Number of queries
---------------	----------------	----------------	-------------------

	1	2	3
	5	3	20
	15	3	65
	60	3	300

Non Authoritative	Interval (min)	Latency (usec)	Number of queries
-------------------	----------------	----------------	-------------------

	1	2	2
	5	3	10
	15	3	55
	60	3	150

### When network monitoring for DNS is off

```
Infoblox > show monitor
```

```
Network Monitoring for DNS is OFF
```

## show monitor dns alert

The `show monitor dns alert` command displays the current DNS alert thresholds. The appliance displays the default thresholds (50% for both invalid ports and invalid TXIDs) if you have not configured new thresholds for the DNS alerts.

## Syntax

```
show monitor dns alert
```

This command has no arguments.

## Example

### Viewing DNS alert thresholds

```
Infoblox > show monitor dns alert
```

```
DNS Network Monitoring is enabled.
```

```
Alerting is enabled.
```

```
DNS Alert      Threshold (per minute)
=====
port           over 70% of packets
txid          over 100 packets
```

### show monitor dns alert status

The `show monitor dns alert status` command displays the current status of invalid DNS responses that arrive on DNS ports that are not open and have mismatched TXIDs (DNS transaction ID). You can view the alert status to identify the primary source of invalid DNS responses. The appliance displays historical alert counts and up to five primary sources that generate invalid DNS responses.

## Syntax

```
show monitor dns alert status
```

This command has no arguments.

## Example

### Viewing DNS alert status

```
Infoblox > show monitor dns alert status
```

```
Data last updated: Mon Oct 6 14:47:12 2008
```

```
DNS Alert   1m   5m   15m   60m   24h   Ever
=====
```

```
port        8    12   12    12    12    12
txid        8    12   12    12    12    12
```

```
There were 80 DNS responses seen in the last minute.
10% were to an invalid port.
```

```
10% had an invalid TXID.
```

```
Primary sources of invalid responses:
```

```
4.4.4.4 (unknown) sent 4
```

```
2.2.2.2 (unknown) sent 3
```

```
7.7.7.7 (unknown) sent 1
```

## show ms\_sticky\_ip

The `show ms_sticky_ip` command displays whether the `ms_sticky_ip` is turned on or off.

### Syntax

```
show ms_sticky_ip
```

This command has no arguments.

### Example

```
Infoblox > show ms_sticky_ip
```

```
show ms_sticky_ip
```

```
ms_sticky_ip is off
```

## show named\_recv\_sock\_buf\_size

The `show named_recv_sock_buf_size` command displays the current BIND receive socket buffer size. The default is 1,536 kilobytes. For information about how to set the receive socket buffer size, see [set named\\_recv\\_sock\\_buf\\_size](#).

### Syntax

```
show named_recv_sock_buf_size
```

This command has no arguments.

### Example

```
Infoblox > show named_recv_sock_buf_size
```

```
DNS 'named' UDP receive socket buffer size: 5000
```



## show network

The **show network** command displays the current network settings for the NIOS appliance and status with respect to a Grid. For information on how to change your network settings, see [set network](#).

### Syntax

```
show network
```

This command has no arguments.

### Example

```
Infoblox > show network
Current LAN1 network settings:
IP Address: 10.34.33.11
Network Mask: 255.255.255.0
Gateway Address: 10.34.33.1
VLAN Tag: 110
DSCP Value: 23
IPv6 Address: 2620:010A:6000:2400:0000:0000:0000:6508/64
IPv6 Gateway Address: 2620:010A:6000:2400:0000:0000:0000:0001
IPv6 VLAN Tag: Untagged
IPv6 DSCP Value: Inherited

HA enabled: false
Grid Status: Member of Infoblox Grid

Current LAN2 Port Settings:
LAN2 Port enabled: true
NIC failover for LAN1 and LAN2 enabled: false
LAN2 IP Address: 10.1.1.35
LAN2 Netmask: 255.255.255.0
LAN2 Gateway: 10.1.1.1
```

## show ntp

The

```
show ntp
```

command displays a list of the peers of the NTP server, along with status information about each peer.

## Syntax

```
show ntp
```

This command has no arguments.

## Example

```
Infoblox > show ntp
```

```
remote      refid      st t   when   poll  reach  delay  offset  jitter
=====
*LOCAL(1)   LOCAL(1)  12 l   47     64    377   +0.000 0.000  0.008
```

When you execute the `show ntp` command, the NIOS appliance displays the following information:

- **remote**: The IP address of the remote peer.
- **refid**: Identifies the reference clock.
- **st**: The stratum of the remote peer.
- **t**: The type of the peer, such as local, unicast or broadcast.
- **when**: When the last packet was received.
- **poll**: The polling interval, in seconds.
- **reach**: The reachability register, in octal numerals.
- **delay**: The current estimated delay, in seconds.
- **offset**: The offset of the peer clock relative to the local clock, in milliseconds.
- **jitter**: The estimated time error of the system clock.

The following special characters may be seen on a peer:

- **\***: The NTP server is synchronized to this peer.
- **#**: The NTP server is almost synchronized to this peer.
- **+**: This peer is selected for possible synchronization.
- **-**: This peer is a candidate for selection.
- **~**: This Peer is statically configured.

## show ospf

The `show ospf` command displays configuration and statistical information about the OSPF protocol that is running on the NIOS appliance. For information on how to change your OSPF settings, see [set ospf](#).

## Syntax

```
show ospf {route | interface | database | neighbor | configuration}
```

The `show ospf` command displays information about the OSPF configuration on the appliance, reachability information about neighbors, and OSPF routes to destinations. You can specify the command with or without arguments.

Argument	Description
route	Displays the OSPF routing table, as determined by the most recent SPF calculation.
interface	Displays the state and configuration on all interfaces configured with OSPF.
database	Displays all OSPF database information.
neighbor	Displays the OSPF neighbor information.
configuration	Displays the running OSPF configuration file.

## Examples

```
Infoblox > show ospf interface
eth0 is down
  OSPF not enabled on this interface
eth1 is up
  Internet Address 172.32.0.61/24, Broadcast 172.32.0.255, Area 0.0.0.0
  Router ID 172.32.0.61, Network Type BROADCAST, Cost: 100
  Transmit Delay is 1 sec, State DROther, Priority 0
  Designated Router (ID) 172.32.0.110, Interface Address 172.32.0.254
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 3, Retransmit 5
  Hello due in 00:00:01
  Neighbor Count is 1, Adjacent neighbor count is 1
lo is up
  Internet Address 172.16.10.10/32, Area 0.0.0.0
  Router ID 172.32.0.61, Network Type LOOPBACK, Cost: 100
  Transmit Delay is 1 sec, State Loopback, Priority 0
  No designated router on this network
  No backup designated router on this network
```

```
Timer intervals configured, Hello 10, Dead 40, Wait 3, Retransmit 5
Hello due in inactive
Neighbor Count is 0, Adjacent neighbor count is 0
Infoblox > show ospf route
===== OSPF network routing table =====
N 172.16.10.10/32  [100] area: 0.0.0.0
                        directly attached to lo
N 172.32.0.0/24  [100] area: 0.0.0.0
                        directly attached to eth1
===== OSPF router routing table =====
===== OSPF external routing table =====
```

## show overload\_bootp

The **show overload\_bootp** command displays whether the overload BOOTP functionality is turned on or off.

### Syntax

```
show overload_bootp
```

This command has no arguments.

### Example

```
Infoblox > show overload_bootp
Overload BOOTP option turned OFF
```

## show pc\_domain

The **show pc\_domain** command displays the entries of the domains for which the category is overridden. For information about the add category command, see [set pc\\_domain add](#). To see information about the delete category command, see [set pc\\_domain delete](#).

### Syntax

```
show pc_domain
```

This command has no arguments.

## Example

```
Infoblox > show pc_domain
```

```
bbc.com 10308,0,0,0,0,
```

## show phonehome

The `show phonehome` command displays the settings of the phone home feature on the appliance.

## Syntax

```
show phonehome
```

This command has no arguments.

## Example

```
Infoblox > show phonehome
```

```
Phone home enabled <Grid wide>: Yes
```

```
Send to Infoblox Support: Yes
```

```
Support ID: 8561
```

```
Address to send to:
```

```
support@infoblox.com
```

## show query\_capture

The `show query_capture` displays the captured DNS queries that are stored locally on the appliance. You can use filters to restrict the DNS queries to specific time and date range, client FQDNs, and IP addresses. Note that the time zone in the CLI console is the time zone of the Grid member.

---

**Note:** Using multiple CLI commands to filter data for the appliances with large number of captured DNS queries and responses can significantly affect the system performance, protocol performance, and CLI command performance.

---

## Syntax

```
show query_capture tail [<num>]
```

```
follow [[fqdn /regex/] [ip /regex/ | grep /regex/]]
```

```
[[fqdn /regex/] [ip /regex/ | grep /regex/]] [after
```

[date] time]

[before [date] time]

You can specify the command with or without arguments. If you use the `show query_capture` command without any arguments, it displays queries from the current capture file.

If you use the `show query_capture` command with `follow`, `tail`, `before`, `after`, `fqdn`, and `ip` arguments, you can view queries for a specific data and time range, queries sent from the client FQDNs, and IP addresses.

**Note:** A capture file for logging DNS queries and responses is compressed every 10 minutes or when it reaches 100 MB in size, whichever comes sooner. A new current file is created when a capture file is compressed. If you are running the `show query_capture` command with `follow` or `tail` when a new capture file is created, the CLI console may return an error indicating that `captured-dns-xxxxxx` has become `inaccessible: No such file or directory`. In these situations, you must execute the CLI command again.

Argument	Description
<code>tail &lt;num&gt;</code>	Shows the last num lines from the capture file. If you do not specify a number, only the last ten lines are displayed.
<code>follow</code>	Displays continuously the lines till the end of the capture file as lines are appended. In addition, you can use 'follow' with '/regex' to search for FQDNs, IP addresses, and regular expressions.
<code>before [date] time</code>	Displays captured DNS queries before the specified date time starting with the oldest saved log file. Make sure that you specify the date (optional for today's date) in the day-month-year format (example: 25-oct-2014), and time in hour:minutes:seconds format (example: 10:09:30).
<code>after [date] time</code>	Displays captured DNS queries after the specified date time until the end of the current log file. The CLI console will not display DNS queries when you specify the current date and time or a future date and time. Make sure that you specify the date (optional for today's date) in the day-month-year format (example: 02-aug-2014), and time in hour:minutes:seconds format (example: 10:09:30).
<code>/regex/</code>	Matches the pattern specified in regex and displays the matched lines from the capture file. To insert "/" in the search pattern, use escape sequence \. The search string starts at the first forward slash (/) ends at the second forward slash (/), and ignores the remaining. For information about regular expressions, see Appendix D Regular Expressions, in the <i>Infoblox Administrator Guide</i> .
<code>fqdn /regex/</code>	Displays queries sent from the client FQDN. You can combine this option with 'ip /regex/' as well.
<code>ip /regex/</code>	Displays queries sent from the client IP address. You can combine this option with 'fqdn /regex/' as well.
<code>grep /regex/</code>	Displays queries by the regular expressions.

## Examples

```
Infoblox > show query_capture
```

```
09-Aug-2014 09:55:50.872 client 10.35.1.136#57722: query:  
aaaa80.1852a_95.com IN AAAA + (10.35.3.96)
```

```
09-Aug-2014 09:55:50.872 client 10.35.1.136#57722: query:  
aaaa81.1852a_95.com IN AAAA + (10.35.3.96)
```

```
09-Aug-2014 09:55:50.872 client 10.35.1.136#57722: query:  
aaaa79.1852a_95.com IN AAAA + (10.35.3.96)
```

```
09-Aug-2014 09:56:07.322 client 10.35.1.136#57722: query:  
aaaa82.1852a_95.com IN AAAA + (10.35.3.96)
```

```
09-Aug-2014 09:56:07.322 client 10.35.1.136#57722: query:  
aaaa84.1852a_95.com IN AAAA + (10.35.3.96)
```

```
09-Aug-2014 09:56:07.322 client 10.35.1.136#57722: query:  
aaaa85.1852a_95.com IN AAAA + (10.35.3.96)
```

```
09-Aug-2014 09:56:07.323 client 10.35.1.136#57722: query:  
aaaa86.1852a_95.com IN AAAA + (10.35.3.96)
```

```
09-Aug-2014 09:56:07.323 client 10.35.1.136#57722: query:  
aaaa87.1852a_95.com IN AAAA + (10.35.3.96)
```

```
09-Aug-2014 09:56:07.323 client 10.35.1.136#57722: query:  
aaaa88.1852a_95.com IN AAAA + (10.35.3.96)
```

```
Infoblox > show query_capture before 09-Aug-2014 05:00:00
```

Note that the filtering options require significant CPU resources, which may affect service performance. Infoblox recommends that you use this command only when necessary.

```
05-Aug-2014 13:31:29.057 client 10.35.112.5#40027: query: mx74.2000a_42.com  
IN MX + (10.35.3.148)
```

```
05-Aug-2014 13:31:29.057 client 10.35.112.5#40027: query: mx80.2000a_42.com  
IN MX + (10.35.3.148)
```

```
05-Aug-2014 13:31:29.057 client 10.35.112.5#40027: query: mx81.2000a_42.com  
IN MX + (10.35.3.148)
```

```
05-Aug-2014 13:31:29.057 client 10.35.112.5#40027: query: mx79.2000a_42.com  
IN MX + (10.35.3.148)
```

```
05-Aug-2014 13:31:29.057 client 10.35.112.5#40027: query: mx82.2000a_42.com  
IN MX + (10.35.3.148)
```

```
Infoblox > show query_capture after 11-Aug-2014 09:00:00
```

Note that the filtering options require significant CPU resources, which may affect service performance. Infoblox recommends that you use this command only when necessary.

```
18-Aug-2014 05:44:22.589 client 10.35.3.148#36662: query: 1.0.0.127.in-addr.arpa IN PTR
```

```
+E (10.35.3.148)
```

```
18-Aug-2014 05:45:22.470 client 10.35.3.148#56373: query: 1.0.0.127.in-addr.arpa IN PTR
```

```
+E (10.35.3.148)
```

```
20-Aug-2014 06:45:25.330 client 10.35.3.148#35366: query: hal.2000a_1.com IN SOA +ED (127.0.0.1)
```

```
20-Aug-2014 06:45:25.330 client 10.35.3.148#35366: UDP: query: hal.2000a_1.com IN SOA
```

```
response: NXDOMAIN -AE
```

```
Infoblox > show query_capture fqdn /2000a_1.com/
```

Note that the filtering options require significant CPU resources, which may affect service performance. Infoblox recommends that you use this command only when necessary.

```
20-Aug-2014 06:45:22.352 client 10.35.3.148#35366: query: hal.2000a_1.com IN SOA +ED (127.0.0.1)
```

```
20-Aug-2014 06:45:22.352 client 10.35.3.148#35366: UDP: query: hal.2000a_1.com IN SOA
```

```
response: NXDOMAIN -AE
```

```
20-Aug-2014 06:45:22.352 client 10.35.3.148#35366: query: 2000a_1.com IN SOA +ED (127.0.0.1)
```

```
20-Aug-2014 06:45:25.330 client 10.35.3.148#35366: query: hal.2000a_1.com IN SOA +ED (127.0.0.1)
```

```
Infoblox > show query_capture fqdn /2000a_1.com/ grep /SOA/
```

Note that the filtering options require significant CPU resources, which may affect service performance. Infoblox recommends that you use this command only when necessary.

```
20-Aug-2014 07:54:29.457 client 10.35.3.148#35366: query: hal.2000a_1.com IN SOA +ED (127.0.0.1)
```



```
20-Aug-2014 07:54:29.457 client 10.35.3.148#35366: UDP: query:
hal.2000a_1.com IN SOA
response: NXDOMAIN -AE
20-Aug-2014 07:54:29.457 client 10.35.3.148#35366: query: 2000a_1.com IN
SOA +ED (127.0.0.1)
20-Aug-2014 07:54:29.457 client 10.35.3.148#35366: UDP: query: 2000a_1.com
IN SOA
response: NOERROR -AE 2000a_1.com. 28800 IN SOA infoblox.localdomain.
admin.infoblox.com. 503 10800 3600 2419200 900;
```

```
Infoblox > show query_capture after 18-Aug-2014 05:05:05 before 18-Aug-2014
06:05:05
```

Note that the filtering options require significant CPU resources, which may affect service performance. Infoblox recommends that you use this command only when necessary.

```
18-Aug-2014 05:05:22.342 client 10.35.3.148#50669: query: 1.0.0.127.in-
addr.arpa IN PTR
+E (10.35.3.148)
18-Aug-2014 05:06:22.402 client 10.35.3.148#35534: query: 1.0.0.127.in-
addr.arpa IN PTR
+E (10.35.3.148)
18-Aug-2014 05:07:22.338 client 10.35.3.148#43846: query: 1.0.0.127.in-
addr.arpa IN PTR
+E (10.35.3.148)
18-Aug-2014 05:08:22.498 client 10.35.3.148#50606: query: 1.0.0.127.in-
addr.arpa IN PTR
+E (10.35.3.148)
18-Aug-2014 05:53:22.359 client 10.35.3.148#56078: query: 1.0.0.127.in-
addr.arpa IN PTR
+E (10.35.3.148)
18-Aug-2014 05:54:22.281 client 10.35.3.148#60212: query: 1.0.0.127.in-
addr.arpa IN PTR
+E (10.35.3.148)
18-Aug-2014 05:55:22.368 client 10.35.3.148#35600: query: 1.0.0.127.in-
addr.arpa IN PTR
+E (10.35.3.148)
```

## show remote\_console

The **show remote\_console** command displays remote console access settings. This command allows you to see if remote console access is enabled without accessing the GUI. You can set this option at the member level and the Grid level. For information on how to change your remote\_console, see [set remote\\_console](#).

### Syntax

```
show remote_console
```

This command has no arguments.

### Example

```
Infoblox > show remote_console
current remote console access settings:
Remote console access enabled (Grid-level): true
```

## show reporting\_user\_capabilities

The **show reporting\_user\_capabilities** command displays information about users who have the `delete` permission for reporting data. To set the `delete` permission on reporting data for a local admin who has superuser permissions, see [set reporting\\_user\\_capabilities](#).



#### Note

This command is supported only on the Grid Master.

### Syntax

```
show reporting_user_capabilities
```

### Example

```
Infoblox > show reporting_user_capabilities
User : user1
Capabilities: Delete reporting indexed data

User : user2
Capabilities: Delete reporting indexed data
```

## show restart\_anycast\_with\_dns\_restart

The `show restart_anycast_with_dns_restart` command displays the status of the `set restart_anycast_with_dns_restart` command. For information about how to set restart anycast command, see [set restart\\_anycast\\_with\\_dns\\_restart](#).

### Syntax

```
show restart_anycast_with_dns_restart
```

This command has no arguments.

### Example

The following is an example of the output displayed when you execute the `show restart_anycast_with_dns_restart` command:

```
Infoblox > show restart_anycast_with_dns_restart
```

```
restart_anycast_with_dns_restart is set to "on" (default value).
```

```
Therefore, when the DNS service restarts, the anycast service also restarts.
```

## show routes

The `show routes` command displays the current IPv4 and IPv6 routing information on the NIOS appliance and organizes the information according to the interface. This command is a valuable diagnostic tool for connectivity issues.

If you had selected the **Enable default route redundancy on LAN1/LAN2** checkbox, then output of this command displays two default routes each having a different metric number. For more information about the checkbox, see [About Route Redundancy](#).

### Syntax

```
show routes
```

This command has no arguments.

### Example

In the following example, `default` specifies the direct connection to the interface and the static routes are represented (in this example) in bold. You specify static routes by manually entering them through the GUI.

```
Infoblox > show routes
```

```
From LAN:
```

```
10.34.33.0/24 dev eth1 scope link default via 10.34.33.1 dev eth1
```

```
From LAN2:
```

```
10.1.1.0/24 dev eth3 scope link default via 10.1.1.1 dev eth3
```

```
From IPv4 main route table:
```

```
10.34.33.0/24 dev eth1 proto kernel scope link src 10.34.33.11 10.1.1.0/24  
dev eth3 proto kernel scope link src 10.1.1.35 default via 10.34.33.1 dev  
eth1
```

```
From IPv6 main route table:
```

```
fe80::/64 dev eth1 metric 256 expires 21257697sec mtu 1500 advmss 1440  
metric10 4294967295
```

```
fe80::/64 dev eth3 metric 256 expires 21334065sec mtu 1500 advmss 1440  
metric10 4294967295
```

```
ff00::/8 dev eth1 metric 256 expires 21257697sec mtu 1500 advmss 1440  
metric10 4294967295
```

```
ff00::/8 dev eth3 metric 256 expires 21334065sec mtu 1500 advmss 1440  
metric10 4294967295
```

```
default via fe80::204:96ff:fe1d:1980 dev eth1 proto kernel metric 1024  
expires 1661sec mtu 1500 advmss 1440 metric10 64
```

```
default via fe80::204:96ff:fe1d:1980 dev eth3 proto kernel metric 1024  
expires 1661sec mtu 1500 advmss 1440 metric10 64
```

```
unreachable default dev lo proto none metric -1 error -101 metric10 255
```

The following example shows the output when the **Enable default route redundancy on LAN1/LAN2** checkbox is selected. It displays two default routes each having a different metric number. For IPv4 networks, the primary default route has a metric value of 0 and the secondary default route has a metric value of 10. For IPv6 networks, the primary default route has a metric value of 1024 and the secondary default route has a metric value of 1124.

```
Infoblox > show routes
```

```
From LAN1:
```

```
IPv4 route table:
```

```
10.33.55.0/24 dev oct1 scope link  
default via 10.33.55.1 dev oct1
```

```
IPv6 route table:
```

```
2620:10a:6000:2290::/64 dev oct1 metric 1024
```

```
default via 2620:10a:6000:2290::1 dev oct1 metric 1024
```

```

From IPv4 main route table:
default via 10.34.1.17 dev oct3 proto ifplugd
default via 10.34.52.1 dev oct1 proto ifplugd metric 10
10.33.7.11/29 dev oct3 proto kernel scope link src 10.33.7.11
10.33.7.11/29 dev oct3.103 proto kernel scope link src 10.33.7.11
10.33.55.0/24 dev oct1 proto kernel scope link src 10.33.55.19
10.33.55.0/24 dev oct1.245 proto kernel scope link src 10.33.55.19
10.33.0.0/16 dev oct0 proto kernel scope link src 10.33.1.111
From IPv6 main route table:
2620:10a:6000:2203::/64 dev oct3 proto kernel metric 256
2620:10a:6000:2203::/64 dev oct3.103 proto kernel metric 256
2620:10a:6000:2290::/64 dev oct1 proto kernel metric 256
2620:10a:6000:2290::/64 dev oct1.245 proto kernel metric 256
2620:10a:6000:2500::/64 dev oct0 proto kernel metric 256
fe80::/64 dev oct4 proto kernel metric 256
fe80::/64 dev oct1 proto kernel metric 256
fe80::/64 dev oct1.245 proto kernel metric 256
fe80::/64 dev oct3 proto kernel metric 256
fe80::/64 dev oct3.103 proto kernel metric 256
fe80::/64 dev oct0 proto kernel metric 256
default via 2620:10a:6000:2203::1 dev oct3 proto ifplugd metric 1024
default via 2620:10a:6000:2290::1 dev oct1 proto ifplugd metric 1124

```

## show rpz\_recursive\_only

Use the `show rpz_recursive_only <view_name>` command to view whether NIOS RPZ zones are used instead of local RPZ zones to block records with private IP addresses from being queried by external users. This command is available only on the Gird Master.

For information about changing the setting, run the `set rpz_recursive_only` command.

## Syntax

```
show rpz_recursive_only <view_name> [zone_name]
```

Argument	Description
view_name	DNS view to which the RPZ zones belong.
zone_name	NIOS RPZ zone name to be used

The command can return the following output values:

- `none` : Denotes that the setting applied earlier will continue to apply.
- `yes` : Denotes that NIOS RPZ zones are used instead of local RPZ zones to block records with private IP addresses from being queried by external users.
- `no` : Denotes that local RPZ zones are used to block records with private IP addresses from being queried by external users.

If you do not specify the zone name, the command displays the output for the DNS view only.

## Example

```
Infoblox > show rpz_recursive_only default
default: recursive-only none
```

## show scheduled

Use the `show scheduled` command to view the number of times per hour the appliance checks if the services need a restart when the scheduling task feature is enabled. The appliance restarts services only when the execution of a scheduled task requires a service restart.

Use the `set scheduled` command to configure the value. You can set the value from 0 to 60, and the default value is 60. A value of 0 turns off the restart feature.

## Syntax

```
show scheduled task restarts
```

This command has no argument.

## Example

```
Infoblox > show scheduled task restarts
Number of restarts per hour: 4
```

## show security

The `show security` command shows the current security settings and whether access to the NIOS appliance through the GUI is restricted. For more information, see [set security](#).

## Syntax

```
show security
```

This command has no arguments.

## Example

```
Infoblox > show security
current security settings:
Access restricted: false
```

## show session\_timeout

The `show session_timeout` command shows how long a session remains open when there is no user activity. For more information, see [set session\\_timeout](#).

## Syntax

```
show session_timeout
```

This command has no arguments.

## Example

```
Infoblox > show session_timeout
Current GUI/CLI timeout is 31536000 seconds (8760:00:00)
```

## show smartnic

The `show smartnic` command shows whether monitor mode for the Threat Protection service is on or off. When on, monitor mode for the Threat Protection service is enabled and the appliance logs DNS packets (instead of dropping them) that would have been blocked by threat protection rules. This information is recorded in the audit log. For information on how to disable monitor mode for the Threat Protection service, see [set smartnic monitor-mode](#).

## Syntax

```
show smartnic
```

This command has no arguments.

## Example

```
Infoblox > show smartnic
Firmware version:          3.8.1 Jul 21, 2014,
Log level:                 6
Failed cores:              None
Threat Protection:         Enabled
```

```
Threat Protection monitor mode: Disabled
```

```
Threat Protection event stats: CRITICAL=0 MAJOR=0 WARNING=0
```

```
INFORMATIONAL=575349
```

## show snmp

The **show snmp** command shows information about the SNMP object that you specify. You can enter the SNMP object name or OID. This command is similar to the SNMP "get" operation. You can use the optional **v3** command to get the information using SNMPv3. For information about SNMP, see [Monitoring with SNMP](#) in the *Infoblox NIOS Administrator Guide*.

Use the [set snmptrap](#) command to send SNMP traps to management systems you specify.

## Syntax

```
show snmp variable {name of an SNMP variable, in dotted or symbolic format}  
v3 {snmpuser}
```

Argument	Description
<code>name of an SNMP variable</code>	The name or OID (object ID) of the SNMP object you want to retrieve. For example, you can enter sysName.0 or .1.3.6.1.4.1.2021.11.53.0.
<code>snmpuser</code>	The user name of the SNMPv3 user account. This is optional. If you do not provide a user name, the appliance uses the first SNMPv3 user on the list.

## Examples

### Displaying the host name

```
Infoblox > show snmp variable sysName.0
```

```
SNMPv2-MIB::sysName.0 = STRING: ib-10-34-61-253.infoblox.com
```

### Displaying the CPU temperature

```
Infoblox > show snmp variable .1.3.6.1.4.1.7779.3.1.1.2.1.1.0
```

```
IB-PLATFORMONE-MIB::ibCPUtemperature.0 = STRING: +40.75 C
```

### Displaying the host name using SNMPv3

```
Infoblox > show snmp variable sysName.0 v3 SNMPv3User1
```

```
SNMPv2-MIB::sysName.0 = STRING: ib-10-34-10.42.infoblox.com
```



## show static\_routes

Use the `show static_routes` command to display the current static route configuration on your appliance. To configure static routes, use the `set static_route` command.

You can also use the `show routes` command to view the current IPv4 and IPv6 routing information on the NIOS appliance and how the information is organized according to the interfaces.

## Syntax

`show static_routes [v4|v6]`

Argument	Description
<code>v4   v6</code>	Shows IPv4 or IPv6 static routes. If this is not specified, static routes for both IPv4 and IPv6 are displayed.

## Examples

### Displaying IPv4 Static Routes

Infoblox > `show static_routes v4`

Position	Destination	Gateway
1	1.1.1.1/32	192.168.1.11
2	1.1.1.2/32	192.168.1.12
3	1.1.1.11/32	192.168.1.21
4	1.1.1.3/32	192.168.1.13
5	1.1.1.4/32	192.168.1.14
6	1.1.1.5/32	192.168.1.15
7	1.1.1.6/32	192.168.1.16
8	1.1.1.7/32	192.168.1.17
9	1.1.1.8/32	192.168.1.18

10	1.1.1.9/32	192.168.1.19
11	1.1.1.10/32	192.168.1.20

## Displaying IPv6 Static Routes

Infoblox > `show static_routes v6`

Position	Destination	Gateway
1	1111:2222:3333:4444:5555: 6666:7777:1000/125	1111:2222:3333:4444:5555:6666:7777:9999
2	1111:2222:3333:4444:5555: 6666:7777:2000/125	1111:2222:3333:4444:5555:6666:7777:bbbb

## show status

The

`show status`

command shows the Grid and HA status. For a Grid member, the command also shows the host name and the Grid Master IP address. You can use this command to gather information about the current state of a Grid.

## Syntax

`show status`

This command has no arguments.

## Examples

### Status of Grid Master

Infoblox > `show status`

Grid Status: ID Grid Master

HA Status: Not Configured

## Status of Grid Member

```
Infoblox > show status
Grid Status: ID Grid Member
HA Status: Not Configured
Hostname: member1.infoblox.local
Grid Master IP: 10.64.40.31
```

## Status of HA member

```
Infoblox > show status
Grid Status: ID Grid Member
HA Status: Active
Hostname: member1.infoblox.com
Grid Master IP: 10.35.113.15
```

---

### Note:

- If the Grid member uses IPv6 communication protocol to join a dual mode Grid, then IPv6 address of the Grid Master is displayed.
- If a Grid Master has a Management Interface set up, then the `show status` output of its Grid Member shows the Management interface IP address as the Grid Master IP address.

---

## show subscriber\_secure\_data

If you have configured Infoblox Subscriber Services, the command `show subscriber_secure_data` enables you to view information about the subscriber data cached by the collector member. For information about Infoblox Subscriber Services, see [Infoblox Subscriber Services](#).

## Syntax

```
show subscriber_secure_data
show subscriber_secure_data cache_usage
show subscriber_secure_data persist
show subscriber_secure_data never_proxy
show subscriber_secure_data list_ip_space_desc
```

```
show subscriber_secure_data find <subscriber_id>
```

```
show subscriber_secure_data find <ip_addr> <prefix> <local_id>  
<ip_space_desc>
```

```
show subscriber_secure_data search <regex>
```

Argument	Description
cache_usage	Displays the number of subscriber records in the cache.
persist	Displays the data persistence mode of static subscriber records.
never_proxy	Displays the hexadecimal character of the never_proxy setting.
list_ip_space_desc	Displays the list of all configured IP space discriminators.
find <subscriber_id>	Displays the subscriber record for the matching subscriber client ID. You must specify the subscriber client ID (MSISDN).
find <ip_addr> <prefix> <local_id> > <ip_space_desc>	Displays the subscriber record for the matching subscriber client IP address. Enter the subscriber client IP address. Enter the prefix length. Enter the local ID. Enter the IP space discriminator.
search <regex>	Displays the cached subscriber records for the matching regular expressions.

## Example

### Displaying the subscriber cached entries

```
Infoblox > show subscriber_secure_data
```

```
111.111.111.111/32|IPS:10.35.120.10|ACS:Acct-Session-  
Id=9889732d-34590e08;AN0:User-Name=Helen
```

```
child;NAS:NAS-IP-Address=10.35.120.10;IPS:NAS-IP-  
Address=10.35.120.10;SUB:MSISDN=9988182386;SSP:Subscriber-Secure-  
Policy=0000007f;PCP:Parental-Control-Policy=400000000000000000004000d3;|Sat  
Aug 26 12:55:36 2017
```

```
111.111.111.111/32 |IPS:10.120.252.24|ACS:Acct-Session-  
Id=9944732d-34590e08;AN0:User-Name=IPv6
```

```
only;NAS:NAS-IP-Address=10.120.252.24;IPS:NAS-IP-  
Address=10.120.252.24;SUB:MSISDN=9944182386;PXY:Proxy-  
All=1;PXP:PXY_PRI=0ac4065f;PXS:PXY_SEC=0ac4065f;SSP:Subscriber-Secure-  
Policy=7ff7fffb;PCP:Parental-Control-Policy=020003;|Sat  
Aug 26 12:55:37 2017
```

```
2620:10a:6000:2500::6b02/128|IPS:10.36.120.10|ACS:Acct-Session-  
Id=9979732d-34590e08;AN0:User-Name=Cheap
```

```
NoPC;NAS:NAS-IP-Address=10.26.120.10;IPS:NAS-IP-  
Address=10.26.120.10;SUB:MSISDN=9955182386;|Sat  
Aug 26 12:55:37 2017
```

If the `Proxy-All` setting is set to `1` and the `Subscriber-Secure-Policy` and `Parental-Control-Policy` settings do not block this query, then NIOS proxies these queries to the MSP server for further processing. NIOS first sends the query to the primary MSP server (denoted by the `PXY_PRI` parameter) and if the primary MSP server is not available, NIOS sends the query to the secondary MSP server (denoted by the `PXY_SEC` parameter). If you want all the queries in a specific category to be resolved directly by NIOS without proxying to an MSP server, use the `set subscriber_secure_data never_proxy <category hexadecimal_character>` command. The hexadecimal character represents the category that is not proxied to an MSP server. For more information see [set subscriber\\_secure\\_data never\\_proxy](#).

For example, if you have configured an Internet Off policy for a subscriber during a particular timeframe, then if the `Proxy-All` setting is set to `1`, the MPS server processes and blocks all live video streams by terminating the connections during that timeframe.

If you want all sporting videos except videos about tennis to be blocked, run the `set subscriber_secure_data never_proxy <category hexadecimal_character>` command, where `<category hexadecimal_character>` is the category related to tennis videos. Here, only videos about tennis are allowed to be streamed on NIOS systems; all other videos are blocked.

## Displaying the number of subscriber entries in cache

```
Infoblox > show subscriber_secure_data cache_usage
122798 accounting records in the cache.
```

## Displaying the list of IP Space Discriminators

```
Infoblox > show subscriber_secure_data list_ip_space_desc
NAS-IP-Address=10.2.1.1
NAS-IP-Address=10.36.1.10
NAS-IP-Address=10.120.252.24
NAS-IP-Address=10.36.120.10
```

## Displaying the subscriber entries that matches the specific subscriber ID

```
Infoblox > show subscriber_secure_data find 8089991000
collector service provided info about subscriber (SUB:MSISDN=8089991000)
2620:10a:6000:2500::6b02/128|IPS:10.36.120.10|ACS:Acct-Session-
Id=9979732d-34590e08;AN0:User-Name=Cheap NoPC;NAS:NAS-IP-
Address=10.26.120.10;IPS:NAS-IP-Address=10.26.120.10;SUB:MSISDN=8089991000;|
Sat Aug 26 12:55:37 2017
```

## Displaying the subscriber record for the matching subscriber client IP address

```
Infoblox > show subscriber_secure_data find 10.36.0.151 32 N/A N/A
10.36.0.151/32|LID:N/A|IPS:N/A|FLG:|ACS:Acct-Session-
Id=9999732d-34590e08;NAS:NAS-PORT=1813;EXP:Expire-Profile=Fri Dec 29 09\
\ :08\ :43 2017;PXY:Proxy-All=0;UCP:Unknown-Category-Policy=0;DCP:Dynamic-
Category-Policy=0;SSP:Subscriber-Secure-Policy=ffffffff;PCP:Parental-
Control-
Policy=ffffffffffffffffffffffffffffffffffffffff;PXP:PXY_PRI=04040404;PXS:PXY_SEC=040
40404;SUB:Calling-Station-Id=9956182386;IPA:IP6=2620:10a:6000:2500::c901;|
Mon Dec 24 07:57:07 2018
```

## Displaying the subscriber entries that matches the specified regular expression

```
Infoblox > show subscriber_secure_data search 9889732d-34590e08
10.36.111.1/32|IPS:10.35.120.10|ACS:Acct-Session-
Id=9889732d-34590e08;AN0:User-Name=Helen child;NAS:NAS-IP-
```

```

Address=10.35.120.10;IPS:NAS-IP-
Address=10.35.120.10;SUB:MSISDN=9988182386;SSP:Subscriber-Secure-
Policy=0000007f;PCP:Parental-Control-Policy=400000000000000000004000d3;|Sat
Aug 26 22:57:15 2017
10.36.139.1/32|IPS:10.35.120.10|ACS:Acct-Session-
Id=9889732d-34590e08;AN0:User-Name=Assaf Adult;NAS:NAS-IP-
Address=10.35.120.10;IPS:NAS-IP-
Address=10.35.120.10;SUB:MSISDN=9966182386;SSP:Subscriber-Secure-
Policy=0000061f;|Sat Aug 26 22:57:15 2017
2620:10a:6000:2500::6f01/128|IPS:10.35.120.10|ACS:Acct-Session-
Id=9889732d-34590e08;AN0:User-Name=Helen child;NAS:NAS-IP-
Address=10.35.120.10;IPS:NAS-IP-
Address=10.35.120.10;SUB:MSISDN=9988182386;SSP:Subscriber-Secure-
Policy=0000007f;PCP:Parental-Control-Policy=400000000000000000004000d3;|Sat
Aug 26 22:57:15 2017
2620:10a:6000:2500::8b01/128|IPS:10.35.120.10|ACS:Acct-Session-
Id=9889732d-34590e08;AN0:User-Name=Assaf Adult;NAS:NAS-IP-
Address=10.35.120.10;IPS:NAS-IP-
Address=10.35.120.10;SUB:MSISDN=9966182386;SSP:Subscriber-Secure-
Policy=0000061f;|Sat Aug 26 22:57:15 2017

```

## show subscriber\_secure\_data bypass

If you have configured Infoblox Subscriber Services, the `show subscriber_secure_data bypass` command allows you to view the status of the subscriber data bypass for the member, all members of each site on the entire Grid, or all members of the site. For information about how to set data bypass, see [set subscriber\\_secure\\_data bypass](#). For information about Infoblox Subscriber Services, see [Infoblox Subscriber Services](#).

## Syntax

```
show subscriber_secure_data bypass [ grid | site ] [ "site_name" ]
```

Argument	Description
grid	Display the status of the subscriber service policies bypass for all members of each site on the entire Grid.
site	Display the status of the subscriber service policies bypass for all members of the site.

## Examples

```
Infoblox > show subscriber_secure_data bypass site mobile1
```

```
Site: Site1
```

```
Member: dns1.com
```

```
Subscriber Secure Bypass disabled for the member dns1.com .
```

```
Member: dns2.com
```

```
Subscriber Secure Bypass enabled for the member dns2.com .
```

```
Member: dns3.com
```

```
Subscriber Secure Bypass disabled for the member dns3.com .
```

```
Infoblox> show subscriber_secure_data bypass
```

```
Member: dns3.com
```

```
Subscriber Secure Bypass disabled on the member dns3.com .
```

## show subscriber\_secure\_data garbage\_collect

If you have configured Infoblox Subscriber Services, use the `show subscriber_secure_data garbage_collect` command to display the status of garbage collection for the specific member. For information about how to set the garbage collect command, see [set subscriber\\_secure\\_data garbage\\_collect](#).

## Syntax

```
show subscriber_secure_data garbage_collect
```

This command has no arguments.

## Example

```
Infoblox > show subscriber_secure_data garbage_collect
```

```
This member is configured for garbage collection scheduled at 3 AM everyday.
```



## show subscriber\_secure\_data cache\_usage\_stats

If you have configured Infoblox Subscriber Services, the `show subscriber_secure_data cache_usage_stats` command displays the statistics for subscriber static, provisioned, and non-provisioned devices based on the flag you used in subscriber data.

Non-provisioned devices display the count of all the devices you set without a K flag, provisioned devices display the count of all devices you set with a K flag, and static devices display the count of all the devices you set with an S flag.

### Syntax

```
Infoblox > show subscriber_secure_data cache_usage_stats
```

This command has no arguments.

### Example

```
Infoblox > show subscriber_secure_data
10.120.20.94/32|LID:N/A|IPS:N/A|FLG:BK|ACS:Acct-Session-Id=9999732d-34590e08;NAS:NAS-
PORT=1813;BWI:BWFlag=1;PXP:PXY_PRI=0a23d268;PXS:PXY_SEC=0a23d268;SUB:User-
Name=user0;BL=casino.com,rummy.com,gamble.com,google.com,facebook.com,bbc.com,test1.c
om,test2.com,test3.com,test4.com,test5.com,test6.com,test7.com,test8.com,test9.com;WL
=funpoker.com,cardpairs.com,google.com,cnn.com,bbc.com,test1.com,test2.com,test3.com,
test4.com,test5.com,test6.com,test7.com,test8.com,test15.com,test16.cpm;|Tue Apr 18
07:05:41 2023
10.120.20.93/32|LID:N/A|IPS:N/A|FLG:S|ACS:Acct-Session-Id=28de847acde415ab;NAS:NAS-
PORT=1813;SUB:Calling-Station-Id=1101202028;PXY:Proxy-All=1;PCP:Parental-Control-
Policy=0000000000000000000000000000000000000000000020040;PXP:PXY_PRI=0ac4800d;PXS:PXY_SEC=0ac4800d;|
Thu Apr 20 06:57:22 2023
10.120.20.95/32|LID:N/A|IPS:N/A|FLG:|ACS:Acct-Session-Id=9999732d-34590e08;NAS:NAS-
PORT=1813;PXP:PXY_PRI=0a23d268;PXS:PXY_SEC=0a23d268;SUB:User-Name=acm;|Thu Apr 20 06:5
8:05 2023

Infoblox > show subscriber_secure_data cache_usage_stats
Getting cache usage details may take some time.
Static devices           : 1
Provisioned devices     : 1
Non-provisioned devices : 2
```

## show subscriber\_secure\_data never\_proxy

If you have configured Infoblox Subscriber Services, use the `show subscriber_secure_data never_proxy` command to view the hexadecimal character of the `never_proxy` setting.

To set the hexadecimal value of the `never_proxy` category, see the `set subscriber_secure_data never_proxy` command.

## Syntax

```
show subscriber_secure_data never_proxy
```

## Example

```
Infoblox > show subscriber_secure_data never_proxy
never_proxy category set is 000fffffffffffffffffffffffffffff01
```

## show subscriber\_secure\_data persist

If you have configured Infoblox Subscriber Services, you can use the `show subscriber_secure_data persist` command to view the information about the enabled data persistence mode which allows static records to survive restart. For information about Infoblox Subscriber Services, see [Infoblox Subscriber Services](#).

## Syntax

```
show subscriber_secure_data persist
```

## Example

```
Infoblox > show subscriber_secure_data persist
1.2.3.4/32|IPS:N/A|FLG:S||Mon Oct 30 10:55:50 2017 Persistent Subscriber Secure
Data last modified at Mon Oct 30 15:12:07 2017
```

## show support\_access

The `show support_access` command shows whether the support\_access function is enabled. By default, the support\_access function is disabled. For more information on the support\_access function, see [set support\\_access](#).

## Syntax

```
show support_access
```

This command has no arguments.

## Example

```
Infoblox > show support_access
current support access settings:
  Support access enabled (Grid-level): true
```

## show support\_timeout

The `show support_timeout` command displays the custom timeout value set for the support bundle download. For information about how to set the timeout value, see [set support\\_timeout](#).

## Syntax

```
show support_timeout
```

This command has no arguments.

## Example

```
Infoblox > show support_timeout
Support timeout is 2600
```

## show tcp\_timestamps

The `show tcp_timestamps` command shows whether TCP timestamps are enabled or disabled. If timestamps are enabled, you can view them in the traffic capture file.

Typically, you run the `show tcp_timestamps` command before running the [set tcp\\_timestamps](#) command to determine the status of the TCP timestamps. Based on the status (enabled or disabled), you can run `set tcp_timestamps` to enable or disable the timestamps.

## Syntax

```
show tcp_timestamps
```

## Example

When TCP timestamps are enabled:

```
Infoblox > show tcp_timestamps
TCP timestamps are enabled.
```

When TCP timestamps are disabled:

```
Infoblox > show tcp_timestamps
```

TCP timestamps are disabled.

## show tech-support

The `show tech-support` command displays output for all show commands. It is a labor saving command that allows you to view the information provided by all the `show` commands. Using the log argument allows you to save the output to a log file that is included in the support bundle.

### Syntax

```
show tech-support [log]
```

Argument	Description
<code>log</code>	Saves the output to a log file that is included in the support bundle.

### Example

```
Infoblox > show tech-support  
Current date and time: Thu Aug 24 14:06:01 EDT 2008  
Up time : 19:29  
Version : 4.3r2  
Hardware ID: 4dcef037e91a403fe05e10ecd241  
  
License Type      : Grid Expiration Date : 12/20/2009  
License String   : GgAAADJj2tzLRv8GJ7/Ua4wkRcbnS6Vp5V5RxizS  
License Type     : DNS Expiration Date : Permanent  
  
License String   : EQAAADNj3cWUB/FCZaaFa8JoT5ev  
License Type     : DHCP  
Expiration Date  : Permanent  
  
Enter <return> to continue with More lines or enter q<return> to go back to  
command line  
  
License Type     : Grid Maintenance  
Expiration Date  : 12/20/2009
```

```
License String : HAAAADxo18rNWeMKC6ndKsJpRYqpSelr4xJUin6C6bE=
```

```
License Type : NIOS Maintenance
```

```
Expiration Date : 12/20/2006
```

```
License String : HAAAADlkwcrmRfgfJLXaLsJpRYqpSelr4xJUiiXWseE=
```

```
Version : 4.3r2
```

```
SN: 000100e081277a69
```

```
REVERT version is: N/A
```

```
No upgrade history found for this box.
```

## show temperature

The **show temperature** command displays the temperature.

### Syntax

```
show temperature
```

This command has no arguments.

### Example

```
Infoblox > show temperature
```

```
show temperature
```

```
CPU_TEMP: 25 C
```

## show test\_promote\_master

Use the **show test\_promote\_master** command to view the results of the test promotion of a Grid Master Candidate to Grid Master. You can run a test promotion by running the [set test\\_promote\\_master](#) command.

### Syntax

```
show test_promote_master <ip_address>
```

Argument	Description
<ip_address>	IP address of the Grid Master Candidate to be promoted to Grid Master

## Examples

To see the results of the latest run:

```

Infoblox > show test_promote_master status
[Mon Sep 23 10:37:34 2019] - [10.39.166.9]: Test state is STARTED
a1.com          IRDP port (:2114): Message is sent from GMC [Mon Sep 23
10:37:34 2019]
                VPN port (:1194): Message is sent from GMC [Mon Sep 23
10:37:34 2019]
vm-09-75.com IRDP port (:2114): Connected [Mon Sep 23
10:37:35 2019]
                VPN port (:1194): Connected [Mon Sep 23
10:37:35 2019]

```

To see the test results for a specific Grid Master Candidate:

```

Infoblox > show test_promote_master 10.39.166.9
[Wed Sep 25 05:51:24 2019] - [10.35.191.9]: Test state is FINISHED
a1.com IRDP port (:2114): Message was not received by GMC [Wed Sep 25 05:51:20
2019]
                VPN port (:1194): Message was not received by GMC [Wed Sep 25
05:51:20 2019]
a3.com IRDP port (:2114): Message was not received by GMC [Wed Sep 25 05:51:20
2019]
                VPN port (:1194): Message was not received by GMC [Wed Sep 25
05:51:20 2019]
a2.com IRDP port (:2114): Message was not received by GMC [Wed Sep 25 05:51:20
2019]
                VPN port (:1194): Message was not received by GMC [Wed Sep 25
05:51:20 2019]
a4.com IRDP port (:2114): Message was not received by GMC [Wed Sep 25 05:51:20
2019]
                VPN port (:1194): Message was not received by GMC [Wed Sep 25
05:51:20 2019]

```

```
vm-11-77.com IRDP port (:2114): Connected [Wed Sep 25 05:49:21 2019]
```

```
VPN port (:1194): Connected [Wed Sep 25 05:49:21 2019]
```

## show thresholdtrap

The **show thresholdtrap** command displays the trigger and reset values of the SNMP trap for CPU usage. The CPU usage trap is disabled by default, and the trigger value is set at 100 and reset value at 0. For information about how to configure the trigger and reset values, see [set thresholdtrap](#).

## Syntax

```
show thresholdtrap {type}
```

Argument	Description
type	The type of threshold trap. Enter <b>CpuUsage</b> to display the trigger and reset values of the CPU usage trap. The trap is disabled by default, and the trigger value is set at 100 and reset value at 0.

## Example

```
Infoblox > show thresholdtrap CpuUsage
```

```
Trap type: CpuUsage
```

```
trigger: 80
```

```
reset: 71
```

## show token

The **show token** command displays token settings.

## Syntax

```
show token
```

This command has no arguments.

## Example

```
Infoblox > show token
```

```
show token
```

```
The token is not configured
```

## show traffic\_capture\_status

The `show traffic_capture_status` command displays the status of traffic capture on the NIOS appliance. You can use the `set traffic_capture` command to start or stop the traffic capture on a NIOS appliance. For more information, see [set traffic\\_capture](#).

### Syntax

```
show traffic_capture_status
```

This command has no arguments.

### Example

```
Infoblox > show traffic_capture_status  
Traffic capture is stopped.  
3277072 bytes captured.
```

## show upgrade\_history

The `show upgrade_history` command displays the upgrade history of the NIOS appliance, showing how many times the appliance has been upgraded and the versions for each upgrade. It also shows the revert version—the version that you can go back to—which is the version of the software the appliance was (last) running prior to the most recent upgrade.

### Syntax

```
show upgrade_history
```

This command has no arguments.

### Example

```
Infoblox > show upgrade_history  
REVERT version is: 4.0r1  
[2006/08/14 19:05:48] Upgraded to: 4.0r2-4-06070517
```

## show uptime

The `show uptime` command displays the uptime (hours and minutes) of the NIOS appliance since the last reboot. In a test environment, this command can be used as a metric. In a production environment, this command is of less use since the appliance remains continually functional.



## Syntax

```
show uptime
```

This command has no arguments.

## Example

```
Infoblox > show uptime  
Up time : 19:33
```

## show version

The **show version** command displays the current version of the NIOS software that is installed on the NIOS appliance. You can use this information when performing an upgrade to determine what version of the software to upgrade to.

## Syntax

```
show version
```

This command has no arguments.

## Example

```
Infoblox > show version  
Version : 4.0r2  
SN: 000100e081277a69  
Hotfix : N/A
```

## show vpn\_cert\_dates

Use the **show vpn\_cert\_dates** command to display the start and end dates of the Infoblox appliance certificate. This information is also included in the Support Bundle.

## Syntax

```
show vpn_cert_dates
```

This command has no arguments.

## Example

The following is an example of the command:

```
Infoblox > show vpn_cert_dates
Start Date=Dec 13 11:00:00 2003 GMT
End Date=May 20 11:00:00 2019 GMT
```

## show wins\_forwarding

Use the **show wins\_forwarding** command to display the current configuration for WINS packet forwarding for the Grid or a specific Grid member.

For information about how to configure WINS packet forwarding to Microsoft servers, see [set wins\\_forwarding](#).

## Syntax

```
show wins_forwarding
```

This command has no arguments.

## Examples

Execute the command on the Grid member that inherited the Grid settings

```
Infoblox > show wins_forwarding
Grid level WINS forwarding: enabled
Grid level WINS default server IP: 10.35.0.123
Member level WINS forwarding: Use grid setting
```

Execute the command on the Grid member that overrode the Grid settings

```
Infoblox > show wins_forwarding
Grid level WINS forwarding: Enabled
Grid level WINS default server IP: 10.35.0.123
Member level output interface: LAN2
Member level WINS forwarding: Override grid setting
Member level forwarding: Enabled
Member level WINS server IP: 10.35.0.321
Member level output interface: MGMT
```

Execute the command on the Grid member that overrode the Grid settings and packet forwarding was disabled

```
Infoblox > show wins_forwarding
Grid level WINS forwarding: Enabled
Grid level WINS default server IP: 10.35.0.123
Member level output interface: LAN

Member level WINS forwarding: Override grid setting
Member level forwarding: Disabled
```

## shutdown

The

**shutdown**

command halts the NIOS appliance. The appliance is designed to operate continuously. However, if you want to halt the appliance you can do so with the

**shutdown**

command.



### Note

Once you shutdown the appliance using this command, you must manually bring it back up.

## Syntax

**shutdown**

This command has no arguments.

## Example

The following example uses the **shutdown** command.

```
Infoblox >
```

```
shutdown
```

```
SHUT DOWN THE SYSTEM? [y or n]
```

```
y
```

## snmpget

Fetches the information from a discovered device's SNMP data. You specify the IP address or hostname and the SNMP Object ID (also often referred to as an SNMP variable) or its dotted numeric equivalent as defined in the device MIB.

## Syntax

```
snmpget <hostname or IP address> <SNMP OID>
```

## Example

The following example uses the `snmpget` command, specifying the IP address of a device discovered by NIOS, along with the standard Object ID `sysName.0` to look up the hostname string for a device. You will need the community string or privacy key to fetch the information.

```
Infoblox > snmpget 172.22.53.5 sysName.0  
Enter SNMP Version (1, 2c or 3): 2c  
Enter SNMP community string: *****  
Created directory: /var/lib/net-snmp/cert_indexes  
Created directory: /var/lib/net-snmp/mib_indexes  
SNMPv2-MIB::sysName.0 = STRING: DEVsw08
```

## snmpwalk

Obtain a tree of information from a network device, using automatic SNMP GETNEXT commands. In the NIOS administrative shell version of the `snmpwalk` command, you can specify the SNMP version, the community string, and the desired Root Object ID (OID).

## Syntax

```
snmpwalk <hostname or IP address> <SNMP OID>
```

## Example

The following example lists a partial output from querying the root Object ID for a Cisco Nexus 5K switch (this technique is also useful for looking up other Object IDs within a particular device):

```
Infoblox > snmpwalk 172.22.33.5 1.3  
Enter SNMP Version (1, 2c or 3): 2c
```

```

Enter SNMP community string: *****
SNMPv2-MIB::sysDescr.0 = STRING: Cisco NX-OS(tm) n5000, Software (n5000-
uk9), Version 5.1(3)N2(1b), RELEASE SOFTWARE Copyright (c) 2002-2011 by
Cisco Systems, Inc. Device Manager Version 5.2(1), Compiled 8/31/2012
17:00:00
SNMPv2-MIB::sysObjectID.0 = OID: SNMPv2-SMI::enterprises.9.12.3.1.3.798
DISMAN-EVENT-MIB::sysUpTimeInstance = Timeticks: (165293061) 19 days,
3:08:50.61
SNMPv2-MIB::sysContact.0 = STRING: who@where
SNMPv2-MIB::sysName.0 = STRING: DEVsw03
SNMPv2-MIB::sysLocation.0 = STRING: snmplocation
SNMPv2-MIB::sysServices.0 = INTEGER: 70
SNMPv2-MIB::sysORLastChange.0 = Timeticks: (40) 0:00:00.40
SNMPv2-MIB::sysORID.1 = OID: SNMPv2-MIB::snmpMIB
SNMPv2-MIB::sysORID.2 = OID: SNMP-VIEW-BASED-ACM-MIB::vacmBasicGroup
SNMPv2-MIB::sysORID.3 = OID: SNMP-FRAMEWORK-MIB::snmpFrameworkMIBCompliance
SNMPv2-MIB::sysORID.4 = OID: SNMP-MPD-MIB::snmpMPDCompliance
SNMPv2-MIB::sysORID.5 = OID: SNMP-USER-BASED-SM-MIB::usmMIBCompliance
SNMPv2-MIB::sysORDescr.1 = STRING: The MIB module for SNMPv2 entities
SNMPv2-MIB::sysORDescr.2 = STRING: View-based Access Control Model for
SNMP.
SNMPv2-MIB::sysORDescr.3 = STRING: The SNMP Management Architecture MIB.
...

```

## traceroute

The **traceroute** command displays information on the route IPv4/IPv6 packets. You can use this command to determine the path of an IPv4/IPv6 query. This command provides information on the path packets travel and the time it takes to reach the IPv4/IPv6 destination address.

## Syntax

```
traceroute {hostname | ip_address} [ opt ]
```

Argument	Description
<i>hostname</i>	Fully qualified domain name
<i>ip_address</i>	Valid IPv4/IPv6 address of the host

Argument	Description
opt	<p>Can be any of the following:</p> <ul style="list-style-type: none"> <li>• <code>numerical</code> (specifies to not interpret the IP address as a DNS name)</li> <li>• <code>src_addr</code> (specifies the starting or "from" address)</li> <li>• <code>ICMP</code> (specifies to use <code>ping</code>)</li> <li>• <code>v6</code> (specifies IPv6 hostname)</li> </ul>

## Example

The following example shows you how to use the traceroute command.

```

Infoblox > traceroute 10.1.1.1
traceroute to 10.1.1.1 (10.1.1.1), 30 hops max, 40 byte packets
 1 10.1.1.5 (10.1.1.5) 1.951 ms 1.637 ms 1.734 ms
 2 10.1.1.1 (10.1.1.1) 0.248 ms 0.284 ms 0.239 ms

```

## show rabbitmq\_queues

The `show rabbitmq_queues` command checks the status of the RabbitMQ queue. This command only checks the status of the queue, but does not increase the count.

Use this command mainly as a diagnostic command to be used when you encounter problems related to RabbitMQ. Typically, the queue length must be close to 0; big numbers (in thousands or millions) indicate RabbitMQ issues.

Use the `show rabbitmq_queues` command to check the status of the RabbitMQ queue in the following scenarios:

- When you install or remove the Threat Analytics license
- When Threat Analytics is running
- When you start or stop the threat analytics service
- When a large number of domains is queried
- When the Outbound service is running
- When you add or remove an endpoint
- When an HA failover occurs
- When you add a member to a Grid Master
- During a Grid Master promotion
- When you add or remove a reporting member to or from a Grid
- When you install or remove an RPZ license
- When you install or remove a Threat Protection license
- On an IB-FLEX appliance
- When you perform an Amazon Web Services (AWS) Route 53 synchronization for zones and records
- When you restore or back up the database
- After an auto-synchronization of an ND member
- When you restore a Grid

## Syntax

```
show rabbitmq_queues
```

This command has no arguments.

## Example

```
Infoblox > show rabbitmq_queues
Listing queues....
analytics_dnst_blacklist_queue 0
aws_route53_q 0
sddc_gm 0
sddc_vnode 0
Grid information (OID, host name, virtual IP): 0
infoblox.localdomain "IP of the VM"
Total no. nodes
```

## Using NIOS APIs

Infoblox provides a Perl API (application programming interface) to help facilitate the integration of the Infoblox device into network environments. The Infoblox DMAP API (Data and Management API) is a set of packages delivered with the Infoblox device to install the package. Use the usual Perl module installation tools on your management system to install the package. For Windows systems using the ActiveState build of Perl, the package is called PPM ("Perl Package Manager" or "Programmer's Package Manager"). For UNIX systems, the package is called CPAN, a global archive of Perl resources.

You can make REST API calls using Java version 1.7 and later. For detailed information about these APIs, see the Infoblox API Documentation (in PDF format), at [docs.infoblox.com](https://docs.infoblox.com).



## Reference Information

This section provides reference information in the following sections:

- [Configurations Requiring Service Restart](#)
- [Open Source Copyright and License Statements](#)
- [Product Compliance](#)
- [vNIOS Appliances](#)
- [Guidance Documentation](#)
- [CSV Import Reference](#)
- [Supported Expressions for Search Parameters](#)

## Configurations Requiring Service Restart

This topic includes a list of objects that require service restart after you make configuration changes to them. For more information about restarting services, see [Restarting Services](#).

Object	Service	Action	Comment
Grid Security Properties	DNS	Update	
Member Security Properties	DNS	Update	
Authoritative Zone	DNS	Insert, update, delete	If the "Restart if needed" option is enabled, restart is triggered immediately.
DNS Traffic Control Monitor	DNS	Update	All functions
Network Container	DHCP	Insert, update, delete	
Shared Network	DHCP	Insert, delete	
Blacklist Rule	DNS	Insert, update, delete	
Bulk Host	DNS	Insert, update, delete	
Convert Lease	DHCP	Convert	
Convert IPv6 Lease	DHCP	Convert	
Grid DHCP Failover Association	DHCP	Insert, update, delete	
DHCP Fingerprint Filter	DHCP	Update	
Grid DHCP MAC Filter	DHCP	Update, delete	
DHCP Range	DHCP	Insert, update, delete	

Object	Service	Action	Comment
DNS64 synthesis group	DNS	Update	
NS Group and Authoritative zone	DNS	Changing list of secondaries	
Grid DNS Properties	DNS	Update	
Member DNS Properties	DNS	Update	
RolloverKey	DNS	Rollover	
DNSSEC Sign Zone	DNS	Sign	
Fixed Address	DHCP	Insert, update, delete	
Forward Zone	DNS	Insert, update, delete	
Grid DHCP Properties	DHCP	Update	
Host Record	DNS	Insert, update, delete	
Host Name Rewrite Policy	DHCP	Insert, update	
DNS Traffic Control Load Balanced Domain Name	DNS	Insert, update, delete, restore	
DNS Traffic Control Pool	DNS	Update	
DNS Traffic Control Server	DNS	Update	
DNS Traffic Control Topology Rule	DNS	Insert, update, delete	
IPv6 DHCP Range	DHCP	Insert, update, delete	
IPv6 Fixed Address	DHCP	Insert, update, delete	
IPv6 Network	DHCP	Insert, update, delete	
IPv6 Option Space	DHCP	Update	
IPv6 Shared Network	DHCP	Update	
Member DHCP Properties	DHCP	Insert, update	
NAC Filter	DHCP	Update, delete	

Object	Service	Action	Comment
Network	DHCP	Insert, update, delete	
Network View	DHCP	Update, delete	
NS Group	DNS	Insert, update, delete	
NXDOMAIN Rule	DNS	Insert, update, delete	
DHCP Range Option Filter	DHCP	Insert, update, delete	
DHCP Option Space	DHCP	Update	
Order DHCP Ranges	DHCP	Changing order of DHCP ranges	
DHCP Relay Agent Filter	DHCP	Update, delete	
Resize Network	DHCP	Resize	Accessible from the <b>IPAM</b> tab
Response Policy Zone	DNS	Update	
Roaming Host	DHCP	Insert, update, delete	
Ruleset	DNS	Update, delete	
IPv4/IPv6 Shared Network	DHCP	Update	
Split Network	DHCP	Split	Accessible from the <b>IPAM</b> tab
Join Network	DHCP	Join	Accessible from the <b>IPAM</b> tab
Stub Zone	DNS	Insert, update, delete	
View	DNS	Insert, update, delete	
Import Topology Database	DNS	Import	
Named ACL	DNS	Update	
Grid Properties	DNS	Update	
Member Properties	DNS	Insert, update, delete	
RADIUS Authentication Service	DHCP	Update	

Object	Service	Action	Comment
Grid Reporting Properties	DNS, DHCP	Update	
Reporting Ip Block Group	DNS	Update, delete	
Reporting Member Properties	DNS, DHCP	Update	
Thales Luna HSM Group	DNS	Update	
License	DNS	Delete	Restarts for QRD, RPZ, MSMGMT and REPORTING licenses.

## Open Source Copyright and License Statements

Infoblox has made every attempt to adhere to the guidelines for use and contribution to the open source community. Please report back to Infoblox any suspected violations of the copyrights, use of open source contributions via the distribution of binaries and/or source from Infoblox. It is the intent of Infoblox to comply with the open source rules of use, and comply with the various copyrights found in the distribution of the products from Infoblox.

This appendix contains the copyright notices for the binary-only distribution from Infoblox. Source changes are contributed back to the open source community when the copyright holder states this is desired. As stated by the enclosed copyrights, a copy of open source files used in our binary-only distribution is available from Infoblox. There is a nominal cost to obtain a CD containing the source files, to cover our costs of duplication and distribution. To obtain a copy of the source, contact us via e-mail at [info@infoblox.com](mailto:info@infoblox.com), or call us at 1.408.625.4200. The sections in this appendix include:

- [GNU General Public License](#)
- [PCRE - Perl Compatible Regular Expressions](#)
- [The MIT License for Jansson](#)
- [LibYAML](#)
- [GNU General Public License 2.0 for Suricata, Netfilter and IPTables](#)
- [PSF License Agreement for Python 3.5.0](#)
- [PSF License Agreement for Python 2.6](#)
- [Stix](#)
- [Mixbox](#)
- [Libtaxii](#)
- [Importlib](#)
- [Dateutil](#)
- [Cybox](#)
- [Ehcache](#)
- [MIT License](#)
- [INFO-ZIP](#)
- [The PHP License, version 3.01](#)
- [Net-SNMP](#)
- [The Independent JPEG Group's JPEG software](#)
- [The FreeType Project LICENSE](#)
- [COMMON DEVELOPMENT AND DISTRIBUTION LICENSE \(CDDL\)](#)
- [Distributed Computing Laboratory, Emory University](#)
- [TASM](#)
- [OpenJDK](#)
- [AOP Alliance \(Java/J2EE AOP standards\)](#)

- [Eclipse Public License - v 1.0](#)
- [ECLIPSE SOFTWARE](#)
- [Wietse Venema Copyright](#)
- [ZLIB License](#)
- [VIM License](#)
- [OpenSSL License](#)
- [OpenLDAP License](#)
- [David L. Mills Copyright](#)
- [BSD License](#)
- [MIT Kerberos Copyright](#)
- [Lawrence Berkeley Copyright](#)
- [Ian F. Darwin Copyright](#)
- [Thai Open Source Software Center Copyright](#)
- [Carnegie Mellon University Copyright](#)
- [Julian Seward Copyright](#)
- [ISC DHCP Copyright](#)
- [ISC BIND Copyright](#)
- [perl Artistic License](#)
- [Apache Software License, Version 2.0](#)
- [GNU Lesser General Public License](#)

## GNU General Public License

GNU GENERAL PUBLIC LICENSE

Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

### Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundations software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each authors protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

## GNU GENERAL PUBLIC LICENSE

### TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you". Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Programs source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.  
You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.
2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
  - a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
  - b. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
  - c. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program. In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
  - a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
  - b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,

- c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable. If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program. If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances. It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice. This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.
8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns. Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.
10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

## NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

END OF TERMS AND CONDITIONS

## PCRE - Perl Compatible Regular Expressions

### PCRE2 LICENCE

PCRE2 is a library of functions to support regular expressions whose syntax and semantics are as close as possible to those of the Perl 5 language.

Release 10 of PCRE2 is distributed under the terms of the "BSD" licence, as specified below. The documentation for PCRE2, supplied in the "doc"

directory, is distributed under the same terms as the software itself. The data in the testdata directory is not copyrighted and is in the public domain.

The basic library functions are written in C and are freestanding. Also included in the distribution is a just-in-time compiler that can be used to optimize pattern matching. This is an optional feature that can be omitted when the library is built.

### THE BASIC LIBRARY FUNCTIONS

Written by: Philip Hazel Email local part: ph10

Email domain:cam.ac.uk

University of Cambridge Computing Service, Cambridge, England.

Copyright (c) 1997-2016 University of Cambridge All rights reserved.

### PCRE2 JUST-IN-TIME COMPILATION SUPPORT

Written by: Zoltan Herczeg Email local part: hzmester

Email domain:freemail.hu Copyright(c) 2010-2016 Zoltan Herczeg All rights reserved.

### STACK-LESS JUST-IN-TIME COMPILER

Written by: Zoltan Herczeg Email local part: hzmester

Email domain:freemail.hu Copyright(c) 2009-2016 Zoltan Herczeg All rights reserved.

### THE "BSD" LICENCE

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.



- Neither the name of the University of Cambridge nor the names of any contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

End

## The MIT License for Jansson

Copyright <YEAR> <COPYRIGHT HOLDER>

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

## LibYAML

Copyright (c) 2006 Kirill Simonov

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

## GNU General Public License 2.0 for Suricata, Netfilter and IPTables

GNU General Public License

Version 2, June 1991 Copyright © 1989, 1991 Free Software Foundation, Inc. 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software—to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Lesser General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

## TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.  
You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.
2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
  - a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
  - b. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.

- c. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program. In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
  - a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
  - b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
  - c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable. If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program. If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.  
Each version is given a distinguishing version number. If the Program specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.
10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

#### NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

#### END OF TERMS AND CONDITIONS

##### How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

<one line to give the program's name and a brief idea of what it does.> Copyright (C)  
<year> <name of author>

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA 02110-1301 USA.

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

```
Gnomovision version 69, Copyright (C) year name of author Gnomovision comes with ABSOLUTELY NO WARRANTY;
for details type `show w`. This is free software, and you are welcome to redistribute it under certain conditions; type
`show c` for details.
```

The hypothetical commands ``show w`` and ``show c`` should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than ``show w`` and ``show c``; they could even be mouse-clicks or menu items—whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

```
Yoyodyne, Inc., hereby disclaims all copyright interest in the program `Gnomovision' (which makes passes at compilers)
written by James Hacker.
```

```
<signature of Ty Coon>, 1 April 1989 Ty Coon, President of Vice
```

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Lesser General Public License instead of this License.

## PSF License Agreement for Python 3.5.0

1. This LICENSE AGREEMENT is between the Python Software Foundation ("PSF"), and the Individual or Organization ("Licensee") accessing and otherwise using Python 3.5.0 software in source or binary form and its associated documentation.
2. Subject to the terms and conditions of this License Agreement, PSF hereby grants Licensee a nonexclusive, royalty-free, world-wide license to reproduce, analyze, test, perform and/or display publicly, prepare derivative works, distribute, and otherwise use Python 3.5.0 alone or in any derivative version, provided, however, that PSF's License Agreement and PSF's notice of copyright, i.e., "Copyright © 2001-2015 Python Software Foundation; All Rights Reserved" are retained in Python 3.5.0 alone or in any derivative version prepared by Licensee.
3. In the event Licensee prepares a derivative work that is based on or incorporates Python 3.5.0 or any part thereof, and wants to make the derivative work available to others as provided herein, then Licensee hereby agrees to include in any such work a brief summary of the changes made to Python 3.5.0.
4. PSF is making Python 3.5.0 available to Licensee on an "AS IS" basis. PSF MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED. BY WAY OF EXAMPLE, BUT NOT LIMITATION, PSF MAKES NO AND DISCLAIMS ANY REPRESENTATION OR WARRANTY OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF PYTHON 3.5.0 WILL NOT INFRINGE ANY THIRD PARTY RIGHTS.
5. PSF SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF PYTHON 3.5.0 FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS AS A RESULT OF MODIFYING, DISTRIBUTING, OR OTHERWISE USING PYTHON 3.5.0, OR ANY DERIVATIVE THEREOF, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.
6. This License Agreement will automatically terminate upon a material breach of its terms and conditions.
7. Nothing in this License Agreement shall be deemed to create any relationship of agency, partnership, or joint venture between PSF and Licensee. This License Agreement does not grant permission to use PSF trademarks or trade name in a trademark sense to endorse or promote products or services of Licensee, or any third party.
8. By copying, installing or otherwise using Python 3.5.0, Licensee agrees to be bound by the terms and conditions of this License Agreement.

## PSF License Agreement for Python 2.6

1. This LICENSE AGREEMENT is between the Python Software Foundation ("PSF"), and the Individual or Organization ("Licensee") accessing and otherwise using this software ("Python") in source or binary form and its associated documentation.
2. Subject to the terms and conditions of this License Agreement, PSF hereby grants Licensee a nonexclusive, royalty-free, world-wide license to reproduce, analyze, test, perform and/or display publicly, prepare derivative works, distribute, and otherwise use Python alone or in any derivative version, provided, however, that PSF's License Agreement and PSF's notice of copyright, i.e., "Copyright (c) 2001, 2002, 2003, 2004, 2005, 2006 Python Software Foundation; All Rights Reserved" are retained in Python alone or in any derivative version prepared by Licensee.
3. In the event Licensee prepares a derivative work that is based on or incorporates Python or any part thereof, and wants to make the derivative work available to others as provided herein, then Licensee hereby agrees to include in any such work a brief summary of the changes made to Python.
4. PSF is making Python available to Licensee on an "AS IS" basis. PSF MAKES NO REPRESENTATIONS OR WARRANTIES, EXPRESS OR IMPLIED. BY WAY OF EXAMPLE, BUT NOT LIMITATION, PSF MAKES NO AND DISCLAIMS ANY REPRESENTATION OR WARRANTY OF MERCHANTABILITY OR FITNESS FOR ANY PARTICULAR PURPOSE OR THAT THE USE OF PYTHON WILL NOT INFRINGE ANY THIRD PARTY RIGHTS.
5. PSF SHALL NOT BE LIABLE TO LICENSEE OR ANY OTHER USERS OF PYTHON FOR ANY INCIDENTAL, SPECIAL, OR CONSEQUENTIAL DAMAGES OR LOSS AS A RESULT OF MODIFYING, DISTRIBUTING, OR OTHERWISE USING PYTHON, OR ANY DERIVATIVE THEREOF, EVEN IF ADVISED OF THE POSSIBILITY THEREOF.
6. This License Agreement will automatically terminate upon a material breach of its terms and conditions.
7. Nothing in this License Agreement shall be deemed to create any relationship of agency, partnership, or joint venture between PSF and Licensee. This License Agreement does not grant permission to use PSF trademarks or trade name in a trademark sense to endorse or promote products or services of Licensee, or any third party.
8. By copying, installing or otherwise using Python, Licensee agrees to be bound by the terms and conditions of this License Agreement.

## Stix

Copyright (c) 2015, The MITRE Corporation

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of The MITRE Corporation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## Mixbox

Copyright (c) 2015, The MITRE Corporation

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of The MITRE Corporation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## Libtaxii

Copyright (c) 2013, The MITRE Corporation All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of The MITRE Corporation nor the

names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## Importlib

(No license included in the package, this is in the readme)

Purpose  
=====

This package contains the code from `importlib` as found in Python 2.7. It is provided so that people who wish to use `"importlib.import_module()"` with a version of Python prior to 2.7 or in 3.0 have the function readily available. The code in no way deviates from what can be found in the Python 2.7 standard library.

For documentation, see the `'importlib docs'` for Python 2.7.

## Dateutil

dateutil - Extensions to the standard Python datetime module.

Copyright (c) 2003-2011 - Gustavo Niemeyer <[gustavo@niemeyer.net](mailto:gustavo@niemeyer.net)>

Copyright (c) 2012-2014 - Tomi Pievilainen <[tomi.pievilainen@iki.fi](mailto:tomi.pievilainen@iki.fi)>

Copyright (c) 2014- Yaron de Leeuw <[me@jarondl.net](mailto:me@jarondl.net)>

All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the copyright holder nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## Cybox

Copyright (c) 2015, The MITRE Corporation All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of The MITRE Corporation nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT



(INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## Ehcache

The open source Ehcache project is licensed under the Apache 2.0 License. The text of the license is available below:

Copyright 2003-2010 Terracotta, Inc.

Licensed under the Apache License, Version 2.0 (the "License");

you may not use this file except in compliance with the License.

You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied.

See the License for the specific language governing permissions and limitations under the License.

## MIT License

Copyright (c) 2010 Paul T. McGuire

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions: The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

## INFO-ZIP

This is version 2007-Mar-4 of the Info-ZIP license.

The definitive version of this document should be available at <ftp://ftp.info-zip.org/pub/infozip/license.html> indefinitely and a copy at <http://www.info-zip.org/pub/infozip/license.html>.

Copyright (c) 1990-2007 Info-ZIP. All rights reserved.

For the purposes of this copyright and license, "Info-ZIP" is defined as the following set of individuals:

Mark Adler, John Bush, Karl Davis, Harald Denker, Jean-Michel Dubois, Jean-loup Gailly, Hunter Goatley, Ed Gordon, Ian Gorman, Chris Herborth, Dirk Haase, Greg Hartwig, Robert Heath, Jonathan Hudson, Paul Kienitz, David Kirschbaum, Johnny Lee, Onno van der Linden, Igor Mandrichenko, Steve P. Miller, Sergio Monesi, Keith Owens, George Petrov, Greg Roelofs, Kai Uwe Rommel, Steve Salisbury, Dave Smith, Steven M. Schweda, Christian Spieler, Cosmin Truta, Antoine Verheijen, Paul von Behren, Rich Wales, Mike White.

This software is provided "as is," without warranty of any kind, express or implied. In no event shall Info-ZIP or its contributors be held liable for any direct, indirect, incidental, special or consequential damages arising out of the use of or inability to use this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the above disclaimer and the following restrictions:

1. Redistributions of source code (in whole or in part) must retain the above copyright notice, definition, disclaimer, and this list of conditions.
2. Redistributions in binary form (compiled executables and libraries) must reproduce the above copyright notice, definition, disclaimer, and this list of conditions in documentation and/or other materials provided with the distribution. The sole exception to this condition is redistribution of a standard UnZipSFX binary (including SFXWiz) as part of a self-extracting archive; that is permitted without inclusion of this license, as long as the normal SFX banner has not been removed from the binary or disabled.
3. Altered versions-including, but not limited to, ports to new operating systems, existing ports with new graphical interfaces, versions with modified or added functionality, and dynamic, shared, or static library versions not from Info-ZIP must be plainly marked as such and must not be misrepresented as being the original source or, if binaries, compiled from the original source. Such altered versions also must not be misrepresented as being Info-ZIP releases-including, but not limited to, labeling of the altered versions with the names "Info-ZIP" (or any variation thereof, including, but not limited to, different capitalizations), "Pocket UnZip," "WiZ" or "MacZip" without the explicit permission of Info-ZIP. Such altered versions are further prohibited from misrepresentative use of the Zip-Bugs or Info-ZIP e-mail addresses or the Info-ZIP URL(s), such as to imply Info-ZIP will provide support for the altered versions.
4. Info-ZIP retains the right to use the names "Info-ZIP," "Zip," "UnZip," "UnZipSFX," "WiZ," "Pocket UnZip," "Pocket Zip," and "MacZip" for its own source and binary releases.

## The PHP License, version 3.01

Copyright (c) 1999 - 2006 The PHP Group. All rights reserved. Redistribution and use in source and binary forms, with or without modification, is permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name "PHP" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact [group@php.net](mailto:group@php.net).
4. Products derived from this software may not be called "PHP", nor may "PHP" appear in their name, without prior written permission from [group@php.net](mailto:group@php.net). You may indicate that your software works in conjunction with PHP by saying "Foo for PHP" instead of calling it "PHP Foo" or "phpfoo"
5. The PHP Group may publish revised and/or new versions of the license from time to time. Each version will be given a distinguishing version number.  
Once covered code has been published under a particular version of the license, you may always continue to use it under the terms of that version. You may also choose to use such covered code under the terms of any subsequent version of the license published by the PHP Group. No one other than the PHP Group has the right to modify the terms applicable to covered code created under this License.
6. Redistributions of any form whatsoever must retain the following acknowledgment:  
"This product includes PHP software, freely available from <http://www.php.net/software/>".

THIS SOFTWARE IS PROVIDED BY THE PHP DEVELOPMENT TEAM ``AS IS'' AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE PHP DEVELOPMENT TEAM OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR

SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

-----

This software consists of voluntary contributions made by many individuals on behalf of the PHP Group.

The PHP Group can be contacted via Email at [group@php.net](mailto:group@php.net).

For more information on the PHP Group and the PHP project,

please see [<http://www.php.net>](http://www.php.net).

PHP includes the Zend Engine, freely available at [<http://www.zend.com>](http://www.zend.com).

## Net-SNMP

Various copyrights apply to this package, listed in various separate parts below. Please make sure that you read all the parts. Up until 2001, the project was based at UC Davis, and the first part covers all code written during this time. From 2001 onwards, the project has been based at SourceForge, and Networks Associates Technology, Inc hold the copyright on behalf of the wider Net-SNMP community, covering all derivative work done since then. An additional copyright section has been added as Part 3 below also under a BSD license for the work contributed by Cambridge Broadband Ltd. to the project since 2001.

An additional copyright section has been added as Part 4 below also under a BSD license for the work contributed by Sun Microsystems, Inc. to the project since 2003.

Code has been contributed to this project by many people over

the years it has been in development, and a full list of contributors can be found in the README file under the THANKS section.

---- Part 1: CMU/UCD copyright notice: (BSD like) ----

Copyright 1989, 1991, 1992 by Carnegie Mellon University

Derivative Work - 1996, 1998-2000

Copyright 1996, 1998-2000 The Regents of the University of California

All Rights Reserved

Permission to use, copy, modify and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of CMU and The Regents of the University of California not be used in advertising or publicity pertaining to distribution of the software without specific written permission.

CMU AND THE REGENTS OF THE UNIVERSITY OF CALIFORNIA DISCLAIM ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL CMU OR THE REGENTS OF THE UNIVERSITY OF CALIFORNIA BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM THE LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

---- Part 2: Networks Associates Technology, Inc copyright notice (BSD) ----

Copyright (c) 2001-2003, Networks Associates Technology, Inc  
All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Networks Associates Technology, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 3: Cambridge Broadband Ltd. copyright notice (BSD) ----

Portions of this code are copyright (c) 2001-2003, Cambridge Broadband Ltd. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The name of Cambridge Broadband Ltd. may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 4: Sun Microsystems, Inc. copyright notice (BSD) ----

Copyright © 2003 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054, U.S.A. All rights reserved.

Use is subject to license terms below.

This distribution may include materials developed by third parties.

Sun, Sun Microsystems, the Sun logo and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and other countries.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of the Sun Microsystems, Inc. nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 5: Sparta, Inc copyright notice (BSD) ----

Copyright (c) 2003-2005, Sparta, Inc All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Sparta, Inc nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS ``AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 6: Cisco/BUPTNIC copyright notice (BSD) ----

Copyright (c) 2004, Cisco, Inc and Information Network

Center of Beijing University of Posts and Telecommunications. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- Neither the name of Cisco, Inc, Beijing University of Posts and Telecommunications, nor the names of their contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDERS OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

---- Part 7: Fabasoft R&D Software GmbH & Co KG copyright notice (BSD) ----

Copyright (c) Fabasoft R&D Software GmbH & Co KG, 2003[ [oss@fabasoft.com](mailto:oss@fabasoft.com) |mailto:[oss@fabasoft.com](mailto:oss@fabasoft.com)]

Author: Bernhard Penz <[bernhard.penz@fabasoft.com](mailto:bernhard.penz@fabasoft.com)>

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

- Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
- Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
- The name of Fabasoft R&D Software GmbH & Co KG or any of its subsidiaries, brand or product names may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDER "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT HOLDER BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## The Independent JPEG Group's JPEG software

In plain English:

1. We don't promise that this software works. (But if you find any bugs, please let us know!)

2. You can use this software for whatever you want. You don't have to pay us.
3. You may not pretend that you wrote this software. If you use it in a program, you must acknowledge somewhere in your documentation that you've used the IJG code.

In legalese:

The authors make NO WARRANTY or representation, either express or implied, with respect to this software, its quality, accuracy, merchantability, or fitness for a particular purpose. This software is provided "AS IS", and you, its user, assume the entire risk as to its quality and accuracy.

This software is copyright (C) 1991-1998, Thomas G. Lane. All Rights Reserved except as specified below.

Permission is hereby granted to use, copy, modify, and distribute this software (or portions thereof) for any purpose, without fee, subject to these conditions:

- (1) If any part of the source code for this software is distributed, then this README file must be included, with this copyright and no-warranty notice unaltered; and any additions, deletions, or changes to the original files must be clearly indicated in accompanying documentation.
- (2) If only executable code is distributed, then the accompanying documentation must state that "this software is based in part on the work of the Independent JPEG Group".
- (3) Permission for use of this software is granted only if the user accepts full responsibility for any undesirable consequences; the authors accept NO LIABILITY for damages of any kind.

These conditions apply to any software derived from or based on the IJG code, not just to the unmodified library. If you use our work, you ought to acknowledge us.

Permission is NOT granted for the use of any IJG author's name or company name in advertising or publicity relating to this software or products derived from it. This software may be referred to only as "the Independent JPEG Group's software".

We specifically permit and encourage the use of this software as the basis of commercial products, provided that all warranty or liability claims are assumed by the product vendor.

ansi2knr.c is included in this distribution by permission of L. Peter Deutsch, sole proprietor of its copyright holder, Aladdin Enterprises of Menlo Park, CA. ansi2knr.c is NOT covered by the above copyright and conditions, but instead by the usual distribution terms of the Free Software Foundation; principally, that you must include source code if you redistribute it. (See the file ansi2knr.c for full details.) However, since ansi2knr.c is not needed as part of any program generated from the IJG code, this does not limit you more than the foregoing paragraphs do.

The Unix configuration script "configure" was produced with GNU Autoconf. It is copyright by the Free Software Foundation but is freely distributable. The same holds for its supporting scripts (config.guess, config.sub, ltconfig, ltmain.sh). Another support script, install-sh, is copyright by M.I.T. but is also freely distributable.

It appears that the arithmetic coding option of the JPEG spec is covered by patents owned by IBM, AT&T, and Mitsubishi. Hence arithmetic coding cannot legally be used without obtaining one or more licenses. For this reason, support for arithmetic coding has been removed from the free JPEG software. (Since arithmetic coding provides only a marginal gain over the unpatented Huffman mode, it is unlikely that very many implementations will support it.) So far as we are aware, there are no patent restrictions on the remaining code.

The IJG distribution formerly included code to read and write GIF files. To avoid entanglement with the Unisys LZW patent, GIF reading support has been removed altogether, and the GIF writer has been simplified to produce "uncompressed GIFs". This technique does not use the LZW algorithm; the resulting GIF files are larger than usual, but are readable by all standard GIF decoders.

We are required to state that

"The Graphics Interchange Format(c) is the Copyright property of CompuServe Incorporated. GIF(sm) is a Service Mark property of CompuServe Incorporated."

# The FreeType Project LICENSE

2002-Apr-11

Copyright 1996-2002 by  
David Turner, Robert Wilhelm, and Werner Lemberg

## Introduction

=====

The FreeType Project is distributed in several archive packages; some of them may contain, in addition to the FreeType font engine, various tools and contributions which rely on, or relate to, the FreeType Project.

This license applies to all files found in such packages, and which do not fall under their own explicit license. The license affects thus the FreeType font engine, the test programs, documentation and makefiles, at the very least.

This license was inspired by the BSD, Artistic, and IJG (Independent JPEG Group) licenses, which all encourage inclusion and use of free software in commercial and freeware products alike. As a consequence, its main points are that:

- We don't promise that this software works. However, we will be interested in any kind of bug reports. (‘as is’ distribution)
- You can use this software for whatever you want, in parts or full form, without having to pay us. (‘royalty-free’ usage)
- You may not pretend that you wrote this software. If you use it, or only parts of it, in a program, you must acknowledge somewhere in your documentation that you have used the FreeType code. (‘credits’)

We specifically permit and encourage the inclusion of this software, with or without modifications, in commercial products. We disclaim all warranties covering The FreeType Project and assume no liability related to The FreeType Project.

Finally, many people asked us for a preferred form for a credit/disclaimer to use in compliance with this license. We thus encourage you to use the following text:

"" "

Portions of this software are copyright © 1996-2002 The FreeType Project (www.freetype.org). All rights reserved.

"" "

## Legal Terms

=====

### 0. Definitions

-----

Throughout this license, the terms ‘package’, ‘FreeType Project’, and ‘FreeType archive’ refer to the set of files originally distributed by the authors (David Turner, Robert Wilhelm, and Werner Lemberg) as the ‘FreeType Project’, be they named as alpha, beta or final release.

‘You’ refers to the licensee, or person using the project, where

‘using’ is a generic term including compiling the project’s source code as well as linking it to form a ‘program’ or ‘executable’. This program is referred to as ‘a program using the FreeType engine’.

This license applies to all files distributed in the original FreeType Project, including all source code, binaries and documentation, unless otherwise stated in the file in its original, unmodified form as distributed in the original archive. If you are unsure whether or not a particular file is covered by this license, you must contact us to verify this.

The FreeType Project is copyright (C) 1996-2000 by David Turner, Robert Wilhelm, and Werner Lemberg. All rights reserved except as specified below.



## 1. No Warranty

---

THE FREETYPE PROJECT IS PROVIDED 'AS IS' WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. IN NO EVENT WILL ANY OF THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY DAMAGES CAUSED BY THE USE OR THE INABILITY TO USE, OF THE FREETYPE PROJECT.

## 2. Redistribution

---

This license grants a worldwide, royalty-free, perpetual and irrevocable right and license to use, execute, perform, compile, display, copy, create derivative works of, distribute and sublicense the FreeType Project (in both source and object code forms) and derivative works thereof for any purpose; and to authorize others to exercise some or all of the rights granted herein, subject to the following conditions:

- Redistribution of source code must retain this license file ('FTL.TXT') unaltered; any additions, deletions or changes to the original files must be clearly indicated in accompanying documentation. The copyright notices of the unaltered, original files must be preserved in all copies of source files.
- Redistribution in binary form must provide a disclaimer that states that the software is based in part of the work of the FreeType Team, in the distribution documentation. We also encourage you to put an URL to the FreeType web page in your documentation, though this isn't mandatory.

These conditions apply to any software derived from or based on the FreeType Project, not just the unmodified files. If you use our work, you must acknowledge us. However, no fee need be paid to us.

## 3. Advertising

---

Neither the FreeType authors and contributors nor you shall use the name of the other for commercial, advertising, or promotional purposes without specific prior written permission.

We suggest, but do not require, that you use one or more of the following phrases to refer to this software in your documentation or advertising materials: 'FreeType Project', 'FreeType Engine', 'FreeType library', or 'FreeType Distribution'.

As you have not signed this license, you are not required to accept it. However, as the FreeType Project is copyrighted material, only this license, or another one contracted with the authors, grants you the right to use, distribute, and modify it. Therefore, by using, distributing, or modifying the FreeType Project, you indicate that you understand and accept all the terms of this license.

## 4. Contacts

---

There are two mailing lists related to FreeType:

- [freetype@freetype.org](mailto:freetype@freetype.org)  
Discusses general use and applications of FreeType, as well as future and wanted additions to the library and distribution. If you are looking for support, start in this list if you haven't found anything to help you in the documentation.
- [devel@freetype.org](mailto:devel@freetype.org)  
Discusses bugs, as well as engine internals, design issues, specific licenses, porting, etc.
- <http://www.freetype.org>  
Holds the current FreeType web page, which will allow you to download our latest development version and read online documentation.

You can also contact us individually at:

David Turner <[david.turner@freetype.org](mailto:david.turner@freetype.org)>

Robert Wilhelm <[robert.wilhelm@freetype.org](mailto:robert.wilhelm@freetype.org)>

Werner Lemberg <[werner.lemberg@freetype.org](mailto:werner.lemberg@freetype.org)>

## COMMON DEVELOPMENT AND DISTRIBUTION LICENSE (CDDL)

Version 1.0

### 1. Definitions.

- 1.1. Contributor means each individual or entity that creates or contributes to the creation of Modifications.
- 1.2. Contributor Version means the combination of the Original Software, prior Modifications used by a Contributor (if any), and the Modifications made by that particular Contributor.
- 1.3. Covered Software means (a) the Original Software, or (b) Modifications, or (c) the combination of files containing Original Software with files containing Modifications, in each case including portions thereof.
- 1.4. Executable means the Covered Software in any form other than Source Code.
- 1.5. Initial Developer means the individual or entity that first makes Original Software available under this License.
- 1.6. Larger Work means a work which combines Covered Software or portions thereof with code not governed by the terms of this License.
- 1.7. License means this document.
- 1.8. Licensable means having the right to grant, to the maximum extent possible, whether at the time of the initial grant or subsequently acquired, any and all of the rights conveyed herein.
- 1.9. Modifications means the Source Code and Executable form of any of the following:
  - A. Any file that results from an addition to, deletion from or modification of the contents of a file containing Original Software or previous Modifications;
  - B. Any new file that contains any part of the Original Software or previous Modification; or
  - C. Any new file that is contributed or otherwise made available under the terms of this License.
- 1.10. Original Software means the Source Code and Executable form of computer software code that is originally released under this License.
- 1.11. Patent Claims means any patent claim(s), now owned or hereafter acquired, including without limitation, method, process, and apparatus claims, in any patent Licensable by grantor.
- 1.12. Source Code means (a) the common form of computer software code in which 1.13. You (or Your) means an individual or a legal entity exercising rights under, and complying with all of the terms of, this License. For legal entities, You includes any entity which controls, is controlled by, or is under common control with You. For purposes of this definition, control means (a) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (b) ownership of more than fifty percent (50%) of the outstanding shares or beneficial ownership of such entity.

### 2. License Grants.

#### 2.1. The Initial Developer Grant.

Conditioned upon Your compliance with Section 3.1 below and subject to third party intellectual property claims, the Initial Developer hereby grants You a world-wide, royalty-free, non-exclusive license:

- (a) under intellectual property rights (other than patent or trademark) Licensable by Initial Developer, to use, reproduce, modify, display, perform, sublicense and distribute the Original Software (or portions thereof), with or without Modifications, and/or as part of a Larger Work; and
- (b) under Patent Claims infringed by the making, using or selling of Original Software, to make, have made, use, practice, sell, and offer for sale, and/or otherwise dispose of the Original Software (or portions thereof).

(c) The licenses granted in Sections 2.1(a) and (b) are effective on the date Initial Developer first distributes or otherwise makes the Original Software available to a third party under the terms of this License.

(d) Notwithstanding Section 2.1(b) above, no patent license is granted: (1) for code that You delete from the Original Software, or (2) for infringements caused by: (i) the modification of the Original Software, or (ii) the combination of the Original Software with other software or devices.

## 2.2. Contributor Grant.

Conditioned upon Your compliance with Section 3.1 below and subject to third party intellectual property claims, each Contributor hereby grants You a world-wide, royalty-free, non-exclusive license:

(a) under intellectual property rights (other than patent or trademark) Licensable by Contributor to use, reproduce, modify, display, perform, sublicense and distribute the Modifications created by such Contributor (or portions thereof), either on an unmodified basis, with other Modifications, as Covered Software and/or as part of a Larger Work; and

(b) under Patent Claims infringed by the making, using, or selling of Modifications made by that Contributor either alone and/or in combination with its Contributor Version (or portions of such combination), to make, use, sell, offer for sale, have made, and/or otherwise dispose of: (1) Modifications made by that Contributor (or portions thereof); and (2) the combination of Modifications made by that Contributor with its Contributor Version (or portions of such combination).

(c) The licenses granted in Sections 2.2(a) and 2.2(b) are effective on the date Contributor first distributes or otherwise makes the Modifications available to a third party.

(d) Notwithstanding Section 2.2(b) above, no patent license is granted: (1) for any code that Contributor has deleted from the Contributor Version; (2) for infringements caused by: (i) third party modifications of Contributor Version, or (ii) the combination of Modifications made by that Contributor with other software (except as part of the Contributor Version) or other devices; or (3) under Patent Claims infringed by Covered Software in the absence of Modifications made by that Contributor.

## 3. Distribution Obligations.

### 3.1. Availability of Source Code.

Any Covered Software that You distribute or otherwise make available in Executable form must also be made available in Source Code form and that Source Code form must be distributed only under the terms of this License.

You must include a copy of this License with every copy of the Source Code form of the Covered Software You distribute or otherwise make available. You must inform recipients of any such Covered Software in Executable form as to how they can obtain such Covered Software in Source Code form in a reasonable manner on or through a medium customarily used for software exchange.

### 3.2. Modifications.

The Modifications that You create or to which You contribute are governed by the terms of this License. You represent that You believe Your Modifications are Your original creation(s) and/or You have sufficient rights to grant the rights conveyed by this License.

### 3.3. Required Notices.

You must include a notice in each of Your Modifications that identifies You as the Contributor of the Modification. You may not remove or alter any copyright, patent or trademark notices contained within the Covered Software, or any notices of licensing or any descriptive text giving attribution to any Contributor or the Initial Developer.

### 3.4. Application of Additional Terms.

You may not offer or impose any terms on any Covered Software in Source Code form that alters or restricts the applicable version of this License or the recipients rights hereunder. You may choose to offer, and to charge a fee for, warranty, support, indemnity or liability obligations to one or more recipients of Covered Software. However, you may do so only on Your own behalf, and not on behalf of the Initial Developer or any Contributor. You must make it absolutely clear that any such warranty, support, indemnity or liability obligation is offered by You alone, and You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of warranty, support, indemnity or liability terms You offer.

### 3.5. Distribution of Executable Versions.

You may distribute the Executable form of the Covered Software under the terms of this License or under the terms of a license of Your choice, which may contain terms different from this License, provided that You are in compliance with the terms of this License and that the license for the Executable form does not attempt to limit or alter the recipients rights in the Source Code form from the rights set forth in this License. If You distribute the Covered Software in Executable form under a different license, You must make it absolutely clear that any terms which differ from this License are offered by You alone, not by the Initial Developer or Contributor. You hereby agree to indemnify the Initial Developer and every Contributor for any liability incurred by the Initial Developer or such Contributor as a result of any such terms You offer.

### 3.6. Larger Works.

You may create a Larger Work by combining Covered Software with other code not governed by the terms of this License and distribute the Larger Work as a single product. In such a case, You must make sure the requirements of this License are fulfilled for the Covered Software.

## 4. Versions of the License.

### 4.1. New Versions.

Sun Microsystems, Inc. is the initial license steward and may publish revised and/or new versions of this License from time to time. Each version will be given a distinguishing version number. Except as provided in Section 4.3, no one other than the license steward has the right to modify this License.

### 4.2. Effect of New Versions.

You may always continue to use, distribute or otherwise make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. If the Initial Developer includes a notice in the Original Software prohibiting it from being distributed or otherwise made available under any subsequent version of the License, You must distribute and make the Covered Software available under the terms of the version of the License under which You originally received the Covered Software. Otherwise, You may also choose to use, distribute or otherwise make the Covered Software available under the terms of any subsequent version of the License published by the license steward.

### 4.3. Modified Versions.

When You are an Initial Developer and You want to create a new license for Your Original Software, You may create and use a modified version of this License if You: (a) rename the license and remove any references to the name of the license steward (except to note that the license differs from this License); and (b) otherwise make it clear that the license contains terms which differ from this License.

## 5. DISCLAIMER OF WARRANTY.

COVERED SOFTWARE IS PROVIDED UNDER THIS LICENSE ON AN AS IS BASIS, WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, WITHOUT LIMITATION, WARRANTIES THAT THE COVERED SOFTWARE IS FREE OF DEFECTS, MERCHANTABILITY, FIT FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE COVERED SOFTWARE IS WITH YOU. SHOULD ANY COVERED SOFTWARE PROVE DEFECTIVE IN ANY RESPECT, YOU (NOT THE INITIAL DEVELOPER OR ANY OTHER CONTRIBUTOR) ASSUME THE COST OF ANY NECESSARY SERVICING, REPAIR OR CORRECTION. THIS DISCLAIMER OF WARRANTY CONSTITUTES AN ESSENTIAL PART OF THIS LICENSE. NO USE OF ANY COVERED SOFTWARE IS AUTHORIZED HEREUNDER EXCEPT UNDER THIS DISCLAIMER.

## 6. TERMINATION.

6.1. This License and the rights granted hereunder will terminate automatically if You fail to comply with terms herein and fail to cure such breach within 30 days of becoming aware of the breach. Provisions which, by their nature, must remain in effect beyond the termination of this License shall survive.

6.2. If You assert a patent infringement claim (excluding declaratory judgment actions) against Initial Developer or a Contributor (the Initial Developer or Contributor against whom You assert such claim is referred to as Participant) alleging that the Participant Software (meaning the Contributor Version where the Participant is a Contributor or the Original Software where the Participant is the Initial Developer) directly or indirectly infringes any patent, then any and all rights granted directly or indirectly to You by such Participant, the Initial Developer (if the Initial Developer is not the Participant) and all Contributors under Sections 2.1 and/or 2.2 of this License shall, upon 60 days notice from Participant terminate prospectively and automatically at the expiration of such 60 day notice period, unless if within such 60 day period You withdraw Your claim with respect to the Participant Software against such Participant either unilaterally or pursuant to a written agreement with Participant.

6.3. In the event of termination under Sections 6.1 or 6.2 above, all end user licenses that have been validly granted by You or any distributor hereunder prior to termination (excluding licenses granted to You by any distributor) shall survive termination.

#### 7. LIMITATION OF LIABILITY.

UNDER NO CIRCUMSTANCES AND UNDER NO LEGAL THEORY, WHETHER TORT (INCLUDING NEGLIGENCE), CONTRACT, OR OTHERWISE, SHALL YOU, THE INITIAL DEVELOPER, ANY OTHER CONTRIBUTOR, OR ANY DISTRIBUTOR OF COVERED SOFTWARE, OR ANY SUPPLIER OF ANY OF SUCH PARTIES, BE LIABLE TO ANY PERSON FOR ANY INDIRECT, SPECIAL, INCIDENTAL, OR CONSEQUENTIAL DAMAGES OF ANY CHARACTER INCLUDING, WITHOUT LIMITATION, DAMAGES FOR LOST PROFITS, LOSS OF GOODWILL, WORK STOPPAGE, COMPUTER FAILURE OR MALFUNCTION, OR ANY AND ALL OTHER COMMERCIAL DAMAGES OR LOSSES, EVEN IF SUCH PARTY SHALL HAVE BEEN INFORMED OF THE POSSIBILITY OF SUCH DAMAGES. THIS LIMITATION OF LIABILITY SHALL NOT APPLY TO LIABILITY FOR DEATH OR PERSONAL INJURY RESULTING FROM SUCH PARTYS NEGLIGENCE TO THE EXTENT APPLICABLE LAW PROHIBITS SUCH LIMITATION. SOME JURISDICTIONS DO NOT ALLOW THE EXCLUSION OR LIMITATION OF INCIDENTAL OR CONSEQUENTIAL DAMAGES, SO THIS EXCLUSION AND LIMITATION MAY NOT APPLY TO YOU.

#### 8. U.S. GOVERNMENT END USERS.

The Covered Software is a commercial item, as that term is defined in 48 C.F.R. 2.101 (Oct. 1995), consisting of commercial computer software (as that term is defined at 48 C.F.R. 252.227-7014(a)(1)) and commercial computer software documentation as such terms are used in 48 C.F.R. 12.212 (Sept. 1995). Consistent with 48 C.F.R. 12.212 and 48 C.F.R. 227.7202-1 through 227.7202-4 (June 1995), all U.S. Government End Users acquire Covered Software with only those rights set forth herein. This U.S. Government Rights clause is in lieu of, and supersedes, any other FAR, DFAR, or other clause or provision that addresses Government rights in computer software under this License.

#### 9. MISCELLANEOUS.

This License represents the complete agreement concerning subject matter hereof. If any provision of this License is held to be unenforceable, such provision shall be reformed only to the extent necessary to make it enforceable. This License shall be governed by the law of the jurisdiction specified in a notice contained within the Original Software (except to the extent applicable law, if any, provides otherwise), excluding such jurisdictions conflict-of-law provisions. Any litigation relating to this License shall be subject to the jurisdiction of the courts located in the jurisdiction and venue specified in a notice contained within the Original Software, with the losing party responsible for costs, including, without limitation, court costs and reasonable attorneys fees and expenses. The application of the United Nations Convention on Contracts for the International Sale of Goods is expressly excluded. Any law or regulation which provides that the language of a contract shall be construed against the drafter shall not apply to this License. You agree that You alone are responsible for compliance with the United States export administration regulations (and the export control laws and regulation of any other countries) when You use, distribute or otherwise make available any Covered Software.

#### 10. RESPONSIBILITY FOR CLAIMS.

As between Initial Developer and the Contributors, each party is responsible for claims and damages arising, directly or indirectly, out of its utilization of rights under this License and You agree to work with Initial Developer and Contributors to distribute such responsibility on an equitable basis. Nothing herein is intended or shall be deemed to constitute any admission of liability.

## Distributed Computing Laboratory, Emory University

This software is released to the public domain, in the spirit of the original code written by Doug Lea. The code can be used for any purpose, modified, and redistributed without acknowledgment. No warranty is provided, either express or implied.

## TASM

Copyright (c) 2000-2005 INRIA, France Telecom All rights reserved.  
Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of the copyright holders nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE COPYRIGHT HOLDERS AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE COPYRIGHT OWNER OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## OpenJDK

GNU General Public License, version 2, with the Classpath Exception

The GNU General Public License (GPL) Version 2, June 1991

Copyright (C) 1989, 1991 Free Software Foundation, Inc.  
59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

### Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public License is intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users. This General Public License applies to most of the Free Software Foundation's software and to any other program whose authors commit to using it. (Some other Free Software Foundation software is covered by the GNU Library General Public License instead.) You can apply it to your programs, too.

When we speak of free software, we are referring to freedom, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish), that you receive source code or can get it if you want it, that you can change the software or use pieces of it in new free programs; and that you know you can do these things.

To protect your rights, we need to make restrictions that forbid anyone to deny you these rights or to ask you to surrender the rights. These restrictions translate to certain responsibilities for you if you distribute copies of the software, or if you modify it.

For example, if you distribute copies of such a program, whether gratis or for a fee, you must give the recipients all the rights that you have. You must make sure that they, too, receive or can get the source code. And you must show them these terms so they know their rights.

We protect your rights with two steps: (1) copyright the software, and (2) offer you this license which gives you legal permission to copy, distribute and/or modify the software.

Also, for each author's protection and ours, we want to make certain that everyone understands that there is no warranty for this free software. If the software is modified by someone else and passed on, we want its recipients to know that what they have is not the original, so that any problems introduced by others will not reflect on the original authors' reputations.

Finally, any free program is threatened constantly by software patents. We wish to avoid the danger that redistributors of a free program will individually obtain patent licenses, in effect making the program proprietary. To prevent this, we have made it clear that any patent must be licensed for everyone's free use or not licensed at all.

The precise terms and conditions for copying, distribution and modification follow.

### TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License applies to any program or other work which contains a notice placed by the copyright holder saying it may be distributed under the terms of this General Public License. The "Program", below, refers to any such program or work, and a "work based on the Program" means either the Program or any derivative work under copyright law: that is to say, a work containing the Program or a portion of it, either verbatim or with modifications and/or translated into another language. (Hereinafter, translation is included without limitation in the term "modification".) Each licensee is addressed as "you".

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running the Program is not restricted, and the output from the Program is covered only if its contents constitute a work based on the Program (independent of having been made by running the Program). Whether that is true depends on what the Program does.

1. You may copy and distribute verbatim copies of the Program's source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and give any other recipients of the Program a copy of this License along with the Program.

You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.

2. You may modify your copy or copies of the Program or any portion of it, thus forming a work based on the Program, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions:
  - a. You must cause the modified files to carry prominent notices stating that you changed the files and the date of any change.
  - b. You must cause any work that you distribute or publish, that in whole or in part contains or is derived from the Program or any part thereof, to be licensed as a whole at no charge to all third parties under the terms of this License.
  - c. If the modified program normally reads commands interactively when run, you must cause it, when started running for such interactive use in the most ordinary way, to print or display an announcement including an appropriate copyright notice and a notice that there is no warranty (or else, saying that you provide a warranty) and that users may redistribute the program under these conditions, and telling the user how to view a copy of this License. (Exception: if the Program itself is interactive but does not normally print such an announcement, your work based on the Program is not required to print an announcement.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Program, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Program, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Program.

In addition, mere aggregation of another work not based on the Program with the Program (or with a work based on the Program) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may copy and distribute the Program (or a work based on it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you also do one of the following:
  - a. Accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
  - b. Accompany it with a written offer, valid for at least three years, to give any third party, for a charge no more than your cost of physically performing source distribution, a complete machine-readable copy of the corresponding source code, to be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange; or,
  - c. Accompany it with the information you received as to the offer to distribute corresponding source code. (This alternative is allowed only for noncommercial distribution and only if you received the program in object code or executable form with such an offer, in accord with Subsection b above.)

The source code for a work means the preferred form of the work for making modifications to it. For an executable work, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the executable. However, as a special exception, the source code distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

If distribution of executable or object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place counts as distribution of the source code, even though third parties are not compelled to copy the source along with the object code.

4. You may not copy, modify, sublicense, or distribute the Program except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense or distribute the Program is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
5. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Program or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Program (or any work based on the Program), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Program or works based on it.
6. Each time you redistribute the Program (or any work based on the Program), the recipient automatically receives a license from the original licensor to copy, distribute or modify the Program subject to these terms and conditions. You may not impose any further restrictions on the recipients' exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties to this License.
7. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Program at all. For example, if a patent license would not permit royalty-free redistribution of the Program by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Program.

If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the section is intended to apply and the section as a whole is intended to apply in other circumstances.

It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system, which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

8. If the distribution and/or use of the Program is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Program under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
9. The Free Software Foundation may publish revised and/or new versions of the General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.

Each version is given a distinguishing version number. If the Program specifies a version number of this License



which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Program does not specify a version number of this License, you may choose any version ever published by the Free Software Foundation.

10. If you wish to incorporate parts of the Program into other free programs whose distribution conditions are different, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

#### NO WARRANTY

11. BECAUSE THE PROGRAM IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE PROGRAM, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE PROGRAM "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE PROGRAM IS WITH YOU. SHOULD THE PROGRAM PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
12. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE PROGRAM AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE PROGRAM (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE PROGRAM TO OPERATE WITH ANY OTHER PROGRAMS), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

#### END OF TERMS AND CONDITIONS

#### How to Apply These Terms to Your New Programs

If you develop a new program, and you want it to be of the greatest possible use to the public, the best way to achieve this is to make it free software which everyone can redistribute and change under these terms.

To do so, attach the following notices to the program. It is safest to attach them to the start of each source file to most effectively convey the exclusion of warranty; and each file should have at least the "copyright" line and a pointer to where the full notice is found.

One line to give the program's name and a brief idea of what it does.

Copyright (C) <year> <name of author>

This program is free software; you can redistribute it and/or modify it under the terms of the GNU General Public License as published by the Free Software Foundation; either version 2 of the License, or (at your option) any later version.

This program is distributed in the hope that it will be useful, but WITHOUT ANY WARRANTY; without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the GNU General Public License for more details.

You should have received a copy of the GNU General Public License along with this program; if not, write to the Free Software Foundation, Inc., 59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Also add information on how to contact you by electronic and paper mail.

If the program is interactive, make it output a short notice like this when it starts in an interactive mode:

Gnomovision version 69, Copyright (C) year name of author Gnomovision comes with ABSOLUTELY NO WARRANTY; for details type 'show w'. This is free software, and you are welcome to redistribute it under certain conditions; type 'show c' for details.

The hypothetical commands 'show w' and 'show c' should show the appropriate parts of the General Public License. Of course, the commands you use may be called something other than 'show w' and 'show c'; they could even be mouse-clicks or menu items--whatever suits your program.

You should also get your employer (if you work as a programmer) or your school, if any, to sign a "copyright disclaimer" for the program, if necessary. Here is a sample; alter the names:

Yoyodyne, Inc., hereby disclaims all copyright interest in the program 'Gnomovision' (which makes passes at compilers) written by James Hacker.

signature of Ty Coon, 1 April 1989

Ty Coon, President of Vice

This General Public License does not permit incorporating your program into proprietary programs. If your program is a subroutine library, you may consider it more useful to permit linking proprietary applications with the library. If this is what you want to do, use the GNU Library General Public License instead of this License.

#### "CLASSPATH" EXCEPTION TO THE GPL

Certain source files distributed by Oracle America and/or its affiliates are subject to the following clarification and special exception to the GPL, but only where Oracle has expressly included in the particular source file's header the words "Oracle designates this particular file as subject to the "Classpath" exception as provided by Oracle in the LICENSE file that accompanied this code."

Linking this library statically or dynamically with other modules is making a combined work based on this library. Thus, the terms and conditions of the GNU General Public License cover the whole combination.

As a special exception, the copyright holders of this library give you permission to link this library with independent modules to produce an executable, regardless of the license terms of these independent modules, and to copy and distribute the resulting executable under terms of your choice, provided that you also meet, for each linked independent module, the terms and conditions of the license of that module. An independent module is a module which is not derived from or based on this library. If you modify this library, you may extend this exception to your version of the library, but you are not obligated to do so. If you do not wish to do so, delete this exception statement from your version.

## AOP Alliance (Java/J2EE AOP standards)

LICENSE: all the source code provided by AOP Alliance is Public Domain.

## Eclipse Public License - v 1.0

THE ACCOMPANYING PROGRAM IS PROVIDED UNDER THE TERMS OF THIS ECLIPSE PUBLIC LICENSE ("AGREEMENT"). ANY USE, REPRODUCTION OR DISTRIBUTION OF THE PROGRAM CONSTITUTES RECIPIENT'S ACCEPTANCE OF THIS AGREEMENT.

### 1. DEFINITIONS

"Contribution" means:

- a) in the case of the initial Contributor, the initial code and documentation distributed under this Agreement, and
- b) in the case of each subsequent Contributor:
  - i) changes to the Program, and
  - ii) additions to the Program;

where such changes and/or additions to the Program originate from and are distributed by that particular Contributor. A Contribution 'originates' from a Contributor if it was added to the Program by such Contributor itself or anyone acting on such Contributor's behalf. Contributions do not include additions to the Program which: (i) are separate modules of software distributed in conjunction with the Program under their own license agreement, and (ii) are not derivative works of the Program.

"Contributor" means any person or entity that distributes the Program.

"Licensed Patents " mean patent claims licensable by a Contributor which are necessarily infringed by the use or sale of its Contribution alone or when combined with the Program.

"Program" means the Contributions distributed in accordance with this Agreement.

"Recipient" means anyone who receives the Program under this Agreement, including all Contributors.

## 2. GRANT OF RIGHTS

a) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free copyright license to reproduce, prepare derivative works of, publicly display, publicly perform, distribute and sublicense the Contribution of such Contributor, if any, and such derivative works, in source code and object code form.

b) Subject to the terms of this Agreement, each Contributor hereby grants Recipient a non-exclusive, worldwide, royalty-free patent license under Licensed Patents to make, use, sell, offer to sell, import and otherwise transfer the Contribution of such Contributor, if any, in source code and object code form. This patent license shall apply to the combination of the Contribution and the Program if, at the time the Contribution is added by the Contributor, such addition of the Contribution causes such combination to be covered by the Licensed Patents. The patent license shall not apply to any other combinations which include the Contribution. No hardware per se is licensed hereunder.

c) Recipient understands that although each Contributor grants the licenses to its Contributions set forth herein, no assurances are provided by any Contributor that the Program does not infringe the patent or other intellectual property rights of any other entity. Each Contributor disclaims any liability to Recipient for claims brought by any other entity based on infringement of intellectual property rights or otherwise. As a condition to exercising the rights and licenses granted hereunder, each Recipient hereby assumes sole responsibility to secure any other intellectual property rights needed, if any. For example, if a third party patent license is required to allow Recipient to distribute the Program, it is Recipient's responsibility to acquire that license before distributing the Program.

d) Each Contributor represents that to its knowledge it has sufficient copyright rights in its Contribution, if any, to grant the copyright license set forth in this Agreement.

## 3. REQUIREMENTS

A Contributor may choose to distribute the Program in object code form under its own license agreement, provided that:

a) it complies with the terms and conditions of this Agreement; and

b) its license agreement:

i) effectively disclaims on behalf of all Contributors all warranties and conditions, express and implied, including warranties or conditions of title and non-infringement, and implied warranties or conditions of merchantability and fitness for a particular purpose;

ii) effectively excludes on behalf of all Contributors all liability for damages, including direct, indirect, special, incidental and consequential damages, such as lost profits;

iii) states that any provisions which differ from this Agreement are offered by that Contributor alone and not by any other party; and

iv) states that source code for the Program is available from such Contributor, and informs licensees how to obtain it in a reasonable manner on or through a medium customarily used for software exchange.

When the Program is made available in source code form:

a) it must be made available under this Agreement; and

b) a copy of this Agreement must be included with each copy of the Program.

Contributors may not remove or alter any copyright notices contained within the Program.

Each Contributor must identify itself as the originator of its Contribution, if any, in a manner that reasonably allows subsequent Recipients to identify the originator of the Contribution.

#### 4. COMMERCIAL DISTRIBUTION

Commercial distributors of software may accept certain responsibilities with respect to end users, business partners and the like. While this license is intended to facilitate the commercial use of the Program, the Contributor who includes the Program in a commercial product offering should do so in a manner which does not create potential liability for other Contributors. Therefore, if a Contributor includes the Program in a commercial product offering, such Contributor ("Commercial Contributor") hereby agrees to defend and indemnify every other Contributor ("Indemnified Contributor") against any losses, damages and costs (collectively "Losses") arising from claims, lawsuits and other legal actions brought by a third party against the Indemnified Contributor to the extent caused by the acts or omissions of such Commercial Contributor in connection with its distribution of the Program in a commercial product offering. The obligations in this section do not apply to any claims or Losses relating to any actual or alleged intellectual property infringement. In order to qualify, an Indemnified Contributor must:

a) promptly notify the Commercial Contributor in writing of such claim, and b) allow the Commercial Contributor to control, and cooperate with the Commercial Contributor in, the defense and any related settlement negotiations. The Indemnified Contributor may participate in any such claim at its own expense.

For example, a Contributor might include the Program in a commercial product offering, Product X. That Contributor is then a Commercial Contributor. If that Commercial Contributor then makes performance claims, or offers warranties related to Product X, those performance claims and warranties are such Commercial Contributor's responsibility alone. Under this section, the Commercial Contributor would have to defend claims against the other Contributors related to those performance claims and warranties, and if a court requires any other Contributor to pay any damages as a result, the Commercial Contributor must pay those damages.

#### 5. NO WARRANTY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, THE PROGRAM IS PROVIDED ON AN "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OR CONDITIONS OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Each Recipient is solely responsible for determining the appropriateness of using and distributing the Program and assumes all risks associated with its exercise of rights under this Agreement, including but not limited to the risks and costs of program errors, compliance with applicable laws, damage to or loss of data, programs or equipment, and unavailability or interruption of operations.

#### 6. DISCLAIMER OF LIABILITY

EXCEPT AS EXPRESSLY SET FORTH IN THIS AGREEMENT, NEITHER RECIPIENT NOR ANY CONTRIBUTORS SHALL HAVE ANY LIABILITY FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OR DISTRIBUTION OF THE PROGRAM OR THE EXERCISE OF ANY RIGHTS GRANTED HEREUNDER, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

#### 7. GENERAL

If any provision of this Agreement is invalid or unenforceable under applicable law, it shall not affect the validity or enforceability of the remainder of the terms of this Agreement, and without further action by the parties hereto, such provision shall be reformed to the minimum extent necessary to make such provision valid and enforceable.

If Recipient institutes patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Program itself (excluding combinations of the Program with other software or hardware) infringes such Recipient's patent(s), then such Recipient's rights granted under Section 2(b) shall terminate as of the date such litigation is filed.

All Recipient's rights under this Agreement shall terminate if it fails to comply with any of the material terms or conditions of this Agreement and does not cure such failure in a reasonable period of time after becoming aware of such noncompliance. If all Recipient's rights under this Agreement terminate, Recipient agrees to cease use and distribution of the Program as soon as reasonably practicable. However, Recipient's obligations under this Agreement and any licenses granted by Recipient relating to the Program shall continue and survive.

Everyone is permitted to copy and distribute copies of this Agreement, but in order to avoid inconsistency the Agreement is copyrighted and may only be modified in the following manner. The Agreement Steward reserves the right to publish new versions (including revisions) of this Agreement from time to time. No one other than the Agreement Steward has the right to modify this Agreement. The Eclipse Foundation is the initial Agreement Steward. The Eclipse Foundation may assign the responsibility to serve as the Agreement Steward to a suitable separate entity. Each new version of the Agreement will be given a distinguishing version number. The Program (including Contributions) may always be distributed subject to the version of the Agreement under which it was received. In addition, after a new version of the Agreement is published, Contributor may elect to distribute the Program (including its Contributions) under the new version. Except as expressly stated in Sections 2(a) and 2(b) above, Recipient receives no rights or licenses to the intellectual property of any Contributor under this Agreement, whether expressly, by implication, estoppel or otherwise. All rights in the Program not expressly granted under this Agreement are reserved.

This Agreement is governed by the laws of the State of New York and the intellectual property laws of the United States of America. No party to this Agreement will bring a legal action under this Agreement more than one year after the cause of action arose. Each party waives its rights to a jury trial in any resulting litigation.

#### Related Links

- \* EPL in plain HTML
- \* The EPL on OSI's site
- \* CPL to EPL conversion

## ECLIPSE SOFTWARE

The product includes Eclipse software (the "Eclipse Program") provided by the Eclipse Foundation and licensed to Infoblox Inc. under the Eclipse Public License v1.0.

EXCEPT AS EXPRESSLY SET FORTH IN THE ECLIPSE PUBLIC LICENSE, THE ECLIPSE PROGRAM IS PROVIDED ON AN "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, EITHER EXPRESS OR IMPLIED INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OR CONDITIONS OF TITLE, NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. EXCEPT AS EXPRESSLY SET FORTH IN THE ECLIPSE PUBLIC LICENSE, NEITHER THE ECLIPSE FOUNDATION NOR ANY CONTRIBUTORS TO THE ECLIPSE PROGRAM SHALL HAVE ANY LIABILITY FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING WITHOUT LIMITATION LOST PROFITS), HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THE ECLIPSE PROGRAM, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Any provisions provided by Infoblox relating to the Eclipse Program which differ from the above terms or the Eclipse Public License are offered by Infoblox alone and not by any other party.

The source code for the Eclipse Program is available from Infoblox as described in the open source introduction.

## Wietse Venema Copyright

```
/* ***** * Copyright 1995 by Wietse Venema. All rights reserved.
Some individual * files may be covered by other copyrights. * * This material was originally written and compiled by
Wietse Venema at * Eindhoven University of Technology, The Netherlands, in 1990, 1991, * 1992, 1993, 1994 and 1995.
* * Redistribution and use in source and binary forms are permitted * provided that this entire copyright notice is
duplicated in all such * copies. * * This software is provided "as is" and without any expressed or implied * warranties,
including, without limitation, the implied warranties of * merchantability and fitness for any particular purpose.
***** */
```

## ZLIB License

(C) 1995-2002 Jean-loup Gailly and Mark Adler

This software is provided as-is, without any express or implied warranty. In no event will the authors be held liable for any damages arising from the use of this software.

Permission is granted to anyone to use this software for any purpose, including commercial applications, and to alter it and redistribute it freely, subject to the following restrictions:

The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.

Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software. This notice may not be removed or altered from any source distribution.

Jean-loup Gailly Mark Adler [jloup@gzip.org](mailto:jloup@gzip.org) [madler@alumni.caltech.edu](mailto:madler@alumni.caltech.edu)

If you use the zlib library in a product, we would appreciate **not** receiving lengthy legal documents to sign. The sources are provided for free but without warranty of any kind. The library has been entirely written by Jean-loup Gailly and Mark Adler; it does not include third-party code.

If you redistribute modified sources, we would appreciate that you include in the file ChangeLog history information documenting your changes.

## VIM License

### COPYING

Vim is Charityware. You can use and copy it as much as you like, but you are encouraged to make a donation to orphans in Uganda. Please read the file "runtime/doc/uganda.txt" for details.

There are no restrictions on distributing an unmodified copy of Vim. Parts of Vim may also be distributed, but this text must always be included. You are allowed to include executables that you made from the unmodified Vim sources, your own usage examples and Vim scripts.

If you distribute a modified version of Vim, you are encouraged to send the maintainer a copy, including the source code. Or make it available to the maintainer through ftp; let him know where it can be found. If the number of changes is small (e.g., a modified Makefile) e-mailing the diffs will do. When the maintainer asks for it (in any way) you must make your changes, including source code, available to him.

The maintainer reserves the right to include any changes in the official version of Vim. This is negotiable. You are not allowed to distribute a modified version of Vim when you are not willing to make the source code available to the maintainer.

The current maintainer is Bram Moolenaar <[Bram@vim.org](mailto:Bram@vim.org)>. If this changes, it will be announced in appropriate places (most likely [www.vim.org](http://www.vim.org) and comp.editors). When it is completely impossible to contact the maintainer, the obligation to send him modified source code ceases.

It is not allowed to remove these restrictions from the distribution of the Vim sources or parts of it. These restrictions may also be used for previous Vim releases instead of the text that was included with it.

Vim is Charityware. You can use and copy it as much as you like, but you are encouraged to make a donation for needy children in Uganda. Please see [kcc] below or visit the ICCF web site, available at these mirrors:

<http://iccf-holland.org/> <http://www.vim.org/iccf/> <http://www.iccf.nl/>

The Open Publication License applies to the Vim documentation, see | manual - copyright |.

=== begin of license ===

### VIM LICENSE

There are no restrictions on distributing unmodified copies of Vim except that they must include this license text. You can also distribute unmodified parts of Vim, likewise unrestricted except that they must include this license text. You are also allowed to include executables that you made from the unmodified Vim sources, plus your own usage examples and Vim scripts.

It is allowed to distribute a modified (or extended) version of Vim, including executables and/or source code, when the following four conditions are met:

1. This license text must be included unmodified.
2. The modified Vim must be distributed in one of the following five ways:
  - a. If you make changes to Vim yourself, you must clearly describe in the distribution how to contact you. When the maintainer asks you (in any way) for a copy of the modified Vim you distributed, you must make your changes, including source code, available to the maintainer without fee. The maintainer reserves the right to include your changes in the official version of Vim. What the maintainer will do with your changes and under what license they will be distributed is negotiable. If there has been no negotiation then this license, or a later version, also applies to your changes. The current maintainer is Bram Moolenaar <[Bram@vim.org](mailto:Bram@vim.org)>. If this changes it will be announced in appropriate places (most likely [vim.sf.net](http://vim.sf.net), [www.vim.org](http://www.vim.org) and/or comp.editors). When it is completely impossible to contact the maintainer, the

- obligation to send him your changes ceases. Once the maintainer has confirmed that he has received your changes they will not have to be sent again.
- b. If you have received a modified Vim that was distributed as mentioned under a) you are allowed to further distribute it unmodified, as mentioned at l). If you make additional changes the text under a) applies to those changes.
  - c. Provide all the changes, including source code, with every copy of the modified Vim you distribute. This may be done in the form of a context diff. You can choose what license to use for new code you add. The changes and their license must not restrict others from making their own changes to the official version of Vim.
  - d. When you have a modified Vim which includes changes as mentioned under c), you can distribute it without the source code for the changes if the following three conditions are met:
    - The license that applies to the changes permits you to distribute the changes to the Vim maintainer without fee or restriction, and permits the Vim maintainer to include the changes in the official version of Vim without fee or restriction.
    - You keep the changes for at least three years after last distributing the corresponding modified Vim. When the maintainer or someone who you distributed the modified Vim to asks you (in any way) for the changes within this period, you must make them available to him.
    - You clearly describe in the distribution how to contact you. This contact information must remain valid for at least three years after last distributing the corresponding modified Vim, or as long as possible.
  - e. When the GNU General Public License (GPL) applies to the changes, you can distribute the modified Vim under the GNU GPL version 2 or any later version.
3. A message must be added, at least in the output of the ":version" command and in the intro screen, such that the user of the modified Vim is able to see that it was modified. When distributing as mentioned under 2)e) adding the message is only required for as far as this does not conflict with the license used for the changes.
  4. The contact information as required under 2)a) and 2)d) must not be removed or changed, except that the person himself can make corrections.

If you distribute a modified version of Vim, you are encouraged to use the Vim license for your changes and make them available to the maintainer, including the source code.

The preferred way to do this is by e-mail or by uploading the files to a server and e-mailing the URL. If the number of changes is small (e.g., a modified Makefile) e-mailing a context diff will do. The e-mail address to be used is [<maintainer@vim.org>](mailto:<maintainer@vim.org>)

It is not allowed to remove this license from the distribution of the Vim sources, parts of it or from a modified version. You may use this license for previous Vim releases instead of the license that they came with, at your option.

=== end of license ===

## OpenSSL License

Copyright (c) 1998-2003 The OpenSSL Project. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions

are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit. (<http://www.openssl.org/>)"

4. The names "OpenSSL Toolkit" and "OpenSSL Project" must not be used to endorse or promote products derived from this software without prior written permission. For written permission, please contact [openssl-core@openssl.org](mailto:openssl-core@openssl.org).

5. Products derived from this software may not be called "OpenSSL" nor may "OpenSSL" appear in their names without prior written permission of the OpenSSL Project.

6. Redistributions of any form whatsoever must retain the following acknowledgment:

"This product includes software developed by the OpenSSL Project for use in the OpenSSL Toolkit (<http://www.openssl.org/>)"

THIS SOFTWARE IS PROVIDED BY THE OpenSSL PROJECT "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OpenSSL PROJECT OR ITS CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

This product includes cryptographic software written by Eric Young (eay@cryptsoft.com). This product includes software written by Tim Hudson (tjh@cryptsoft.com).

## OpenLDAP License

The OpenLDAP Public License Version 2.8, 17 August 2003

Redistribution and use of this software and associated documentation ("Software"), with or without modification, are permitted provided that the following conditions are met:

1. Redistributions in source form must retain copyright statements and notices,
2. Redistributions in binary form must reproduce applicable copyright statements and notices, this list of conditions, and the following disclaimer in the documentation and/or other materials provided with the distribution, and
3. Redistributions must contain a verbatim copy of this document.

The OpenLDAP Foundation may revise this license from time to time. Each revision is distinguished by a version number. You may use this Software under terms of this license revision or under the terms of any subsequent revision of the license.

THIS SOFTWARE IS PROVIDED BY THE OPENLDAP FOUNDATION AND ITS CONTRIBUTORS "AS IS" AND ANY EXPRESSED OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE OPENLDAP FOUNDATION, ITS CONTRIBUTORS, OR THE AUTHOR(S) OR OWNER(S) OF THE SOFTWARE BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

The names of the authors and copyright holders must not be used in advertising or otherwise to promote the sale, use or other dealing in this Software without specific, written prior permission. Title to copyright in this Software shall at all times remain with copyright holders.

OpenLDAP is a registered trademark of the OpenLDAP Foundation.

Copyright 1999-2003 The OpenLDAP Foundation, Redwood City, California, USA. All Rights Reserved. Permission to copy and distribute verbatim copies of this document is granted.

## David L. Mills Copyright

Copyright (c) David L. Mills 1992-2003 \*

Permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appears in all copies and that both the copyright notice and this permission notice appear in supporting documentation, and that the name University of Delaware not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. The University



of Delaware makes no representations about the suitability this software for any purpose. It is provided "as is" without express or implied warranty.

## BSD License

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The names of the authors may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

## MIT Kerberos Copyright

Copyright Notice and Legal Administrivia

---

Copyright (C) 1985-2002 by the Massachusetts Institute of Technology. All rights reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. Furthermore if you modify this software you must label your software as modified software and not distribute it in such a fashion that it might be confused with the original MIT software.

M.I.T. makes no representations about the suitability of this software for any purpose. It is provided "as is" without express or implied warranty.

THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

Individual source code files are copyright MIT, Cygnus Support, OpenVision, Oracle, Sun Soft, FundsXpress, and others. Project Athena, Athena, Athena MUSE, Discuss, Hesiod, Kerberos, Moira, and Zephyr are trademarks of the Massachusetts Institute of Technology (MIT). No commercial use of these trademarks may be made without prior written permission of MIT.

"Commercial use" means use of a name in a product or other for-profit manner. It does NOT prevent a commercial firm from referring to the MIT trademarks in order to convey information (although in doing so, recognition of their trademark status should be given).

---

The following copyright and permission notice applies to the OpenVision Kerberos Administration system located in kadmin/create, kadmin/dbutil, kadmin/passwd, kadmin/server, lib/kadm5, and portions of lib/rpc:

Copyright, OpenVision Technologies, Inc., 1996, All Rights Reserved

WARNING: Retrieving the OpenVision Kerberos Administration system source code, as described below, indicates your acceptance of the following terms. If you do not agree to the following terms, do not retrieve the OpenVision Kerberos administration system.

You may freely use and distribute the Source Code and Object Code compiled from it, with or without modification, but

this Source Code is provided to you "AS IS" EXCLUSIVE OF ANY WARRANTY, INCLUDING, WITHOUT LIMITATION, ANY WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE, OR ANY OTHER WARRANTY, WHETHER EXPRESS OR IMPLIED. IN NO EVENT WILL OPENVISION HAVE ANY LIABILITY FOR ANY LOST PROFITS, LOSS OF DATA OR COSTS OF PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES, OR FOR ANY SPECIAL, INDIRECT, OR CONSEQUENTIAL DAMAGES ARISING OUT OF THIS AGREEMENT, INCLUDING, WITHOUT LIMITATION, THOSE RESULTING FROM THE USE OF THE SOURCE CODE, OR THE FAILURE OF THE SOURCE CODE TO PERFORM, OR FOR ANY OTHER REASON.

OpenVision retains all copyrights in the donated Source Code. OpenVision also retains copyright to derivative works of the Source Code, whether created by OpenVision or by a third party. The OpenVision copyright notice must be preserved if derivative works are made based on the donated Source Code.

OpenVision Technologies, Inc. has donated this Kerberos Administration system to MIT for inclusion in the standard Kerberos 5 distribution. This donation underscores our commitment to continuing Kerberos technology development and our gratitude for the valuable work which has been performed by MIT and the Kerberos community.

---

Portions contributed by Matt Crawford <[crawdad@fnal.gov](mailto:crawdad@fnal.gov)> were work performed at Fermi National Accelerator Laboratory, which is operated by Universities Research Association, Inc., under contract DE-AC02-76CHO3000 with the U.S. Department of Energy.

## Lawrence Berkeley Copyright

Copyright (c) 1990 The Regents of the University of California. All rights reserved.

This code is derived from software contributed to Berkeley by Vern Paxson.

The United States Government has rights in this work pursuant to contract no.

DE-AC03-76SF00098 between the United States Department of Energy and the University of California.

Redistribution and use in source and binary forms with or without modification are permitted provided that: (1) source distributions retain this entire copyright notice and comment, and (2) distributions including binaries display the following acknowledgement: "This product includes software developed by the University of California, Berkeley and its contributors" in the documentation or other materials provided with the distribution and in all advertising materials mentioning features or use of this software. Neither the name of the University nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

## Ian F. Darwin Copyright

Copyright (c) Ian F. Darwin 1986, 1987, 1989, 1990, 1991, 1992, 1994, 1995.

Software written by Ian F. Darwin and others; maintained 1994-1999 Christos Zoulas.

This software is not subject to any export provision of the United States Department of Commerce, and may be exported to any country or planet.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice immediately at the beginning of the file, without modification, this list of conditions, and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. All advertising materials mentioning features or use of this software must display the following acknowledgement: This product includes software developed by Ian F. Darwin and others.
4. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR OR

CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## Thai Open Source Software Center Copyright

Copyright (c) 1998, 1999, 2000 Thai Open Source Software Center Ltd and Clark Cooper

Copyright (c) 2001, 2002 Expat maintainers.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, sublicense, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, subject to the following conditions:

The above copyright notice and this permission notice shall be included in all copies or substantial portions of the Software.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

## Carnegie Mellon University Copyright

Copyright (c) 2001 Carnegie Mellon University. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. The name "Carnegie Mellon University" must not be used to endorse or promote products derived from this software without prior written permission. For permission or any other legal details, please contact

Office of Technology Transfer

Carnegie Mellon University

5000 Forbes Avenue

Pittsburgh, PA 15213-3890

(412) 268-4387, fax: (412) 268-7395

[tech-transfer@andrew.cmu.edu](mailto:tech-transfer@andrew.cmu.edu)

4. Redistributions of any form whatsoever must retain the following acknowledgment:

"This product includes software developed by Computing Services at Carnegie Mellon University (<http://www.cmu.edu/computing/>)."

CARNEGIE MELLON UNIVERSITY DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE, INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS, IN NO EVENT SHALL CARNEGIE MELLON UNIVERSITY BE LIABLE FOR ANY SPECIAL, INDIRECT OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF

CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

## Julian Seward Copyright

This program, "bzip2" and associated library "libbzip2", are copyright (C) 1996-2002 Julian R Seward. All rights reserved. Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. The origin of this software must not be misrepresented; you must not claim that you wrote the original software. If you use this software in a product, an acknowledgment in the product documentation would be appreciated but is not required.
3. Altered source versions must be plainly marked as such, and must not be misrepresented as being the original software.
4. The name of the author may not be used to endorse or promote products derived from this software without specific prior written permission.

THIS SOFTWARE IS PROVIDED BY THE AUTHOR "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL THE AUTHOR BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Julian Seward, Cambridge, UK.

[jseward@acm.org](mailto:jseward@acm.org)

bzip2/libbzip2 version 1.0.2 of 30 December 2001

## ISC DHCP Copyright

Copyright (c) 1995, 1996, 1997, 1998, 1999 Internet Software Consortium -DHCP. All rights reserved.

Redistribution and use in source and binary forms, with or without modification, are permitted provided that the following conditions are met:

1. Redistributions of source code must retain the above copyright notice, this list of conditions and the following disclaimer.
2. Redistributions in binary form must reproduce the above copyright notice, this list of conditions and the following disclaimer in the documentation and/or other materials provided with the distribution.
3. Neither the name of Internet Software Consortium - DHCP nor the names of its contributors may be used to endorse or promote products derived from this software without specific prior written permission. THIS SOFTWARE IS PROVIDED BY INTERNET SOFTWARE

CONSORTIUM AND CONTRIBUTORS "AS IS" AND ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE ARE DISCLAIMED. IN NO EVENT SHALL INTERNET SOFTWARE CONSORTIUM OR CONTRIBUTORS BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

## ISC BIND Copyright

Copyright (C) 1996-2002 Internet Software Consortium.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND INTERNET SOFTWARE CONSORTIUM DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL INTERNET SOFTWARE CONSORTIUM BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Portions Copyright (C) 1996-2001 Nomimum, Inc.

Permission to use, copy, modify, and distribute this software for any purpose with or without fee is hereby granted, provided that the above copyright notice and this permission notice appear in all copies.

THE SOFTWARE IS PROVIDED "AS IS" AND NOMINUM DISCLAIMS ALL WARRANTIES WITH REGARD TO THIS SOFTWARE INCLUDING ALL IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS. IN NO EVENT SHALL NOMINUM BE LIABLE FOR ANY SPECIAL, DIRECT, INDIRECT, OR CONSEQUENTIAL DAMAGES OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

## perl Artistic License

The "Artistic License"

### Preamble

The intent of this document is to state the conditions under which a Package may be copied, such that the Copyright Holder maintains some semblance of artistic control over the development of the package, while giving the users of the package the right to use and distribute the Package in a more-or-less customary fashion, plus the right to make reasonable modifications.

### Definitions:

"Package" refers to the collection of files distributed by the Copyright Holder, and derivatives of that collection of files created through textual modification.

"Standard Version" refers to such a Package if it has not been modified, or has been modified in accordance with the wishes of the Copyright Holder as specified below.

"Copyright Holder" is whoever is named in the copyright or copyrights for the package.

"You" is you, if you're thinking about copying or distributing this Package.

"Reasonable copying fee" is whatever you can justify on the basis of media cost, duplication charges, time of people involved, and so on. (You will not be required to justify it to the Copyright Holder, but only to the computing community at large as a market that must bear the fee.)

"Freely Available" means that no fee is charged for the item itself, though there may be fees involved in handling the item. It also means that recipients of the item may redistribute it under the same conditions they received it.

1. You may make and give away verbatim copies of the source form of the Standard Version of this Package without restriction, provided that you duplicate all of the original copyright notices and associated disclaimers.
2. You may apply bug fixes, portability fixes and other modifications derived from the Public Domain or from the Copyright Holder. A Package modified in such a way shall still be considered the Standard Version.
3. You may otherwise modify your copy of this Package in any way, provided that you insert a prominent notice in each changed file stating how and when you changed that file, and provided that you do at least ONE of the following:
  - a. place your modifications in the Public Domain or otherwise make them Freely Available, such as by posting said modifications to Usenet or an equivalent medium, or placing the modifications on a major

- archive site such as uunet.uu.net, or by allowing the Copyright Holder to include your modifications in the Standard Version of the Package.
- b. use the modified Package only within your corporation or organization.
  - c. rename any non-standard executables so the names do not conflict with standard executables, which must also be provided, and provide a separate manual page for each non-standard executable that clearly documents how it differs from the Standard Version.
  - d. make other distribution arrangements with the Copyright Holder.
4. You may distribute the programs of this Package in object code or executable form, provided that you do at least ONE of the following:
    - a. distribute a Standard Version of the executables and library files, together with instructions (in the manual page or equivalent) on where to get the Standard Version.
    - b. accompany the distribution with the machine-readable source of the Package with your modifications.
    - c. give non-standard executables non-standard names, and clearly document the differences in manual pages (or equivalent), together with instructions on where to get the Standard Version.
    - d. make other distribution arrangements with the Copyright Holder.
  5. You may charge a reasonable copying fee for any distribution of this Package. You may charge any fee you choose for support of this Package. You may not charge a fee for this Package itself. However, you may distribute this Package in aggregate with other (possibly commercial) programs as part of a larger (possibly commercial) software distribution provided that you do not advertise this Package as a product of your own. You may embed this Packages interpreter within an executable of yours (by linking); this shall be construed as a mere form of aggregation, provided that the complete Standard Version of the interpreter is so embedded.
  6. The scripts and library files supplied as input to or produced as output from the programs of this Package do not automatically fall under the copyright of this Package, but belong to whoever generated them, and may be sold commercially, and may be aggregated with this Package. If such scripts or library files are aggregated with this Package via the so-called "undump" or "unexec" methods of producing a binary executable image, then distribution of such an image shall neither be construed as a distribution of this Package nor shall it fall under the restrictions of Paragraphs 3 and 4, provided that you do not represent such an executable image as a Standard Version of this Package.
  7. C subroutines (or comparably compiled subroutines in other languages) supplied by you and linked into this Package in order to emulate subroutines and variables of the language defined by this Package shall not be considered part of this Package, but are the equivalent of input as in Paragraph 6, provided these subroutines do not change the language in any way that would cause it to fail the regression tests for the language.
  8. Aggregation of this Package with a commercial distribution is always permitted provided that the use of this Package is embedded; that is, when no overt attempt is made to make this Packages interfaces visible to the end user of the commercial distribution. Such use shall not be construed as a distribution of this Package.
  9. The name of the Copyright Holder may not be used to endorse or promote products derived from this software without specific prior written permission.
  10. THIS PACKAGE IS PROVIDED "AS IS" AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, WITHOUT LIMITATION, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE.

The End

## Apache Software License, Version 2.0

Copyright (c) 2004 The Apache Software Foundation. All rights reserved. TERMS AND CONDITIONS FOR USE, REPRODUCTION, AND DISTRIBUTION

### 1. Definitions.

"License" shall mean the terms and conditions for use, reproduction, and distribution as defined by Sections 1 through 9 of this document.

"Licensor" shall mean the copyright owner or entity authorized by the copyright owner that is granting the License.

"Legal Entity" shall mean the union of the acting entity and all other entities that control, are controlled by, or are under common control with that entity. For the purposes of this definition, "control" means (i) the power, direct or indirect, to cause the direction or management of such entity, whether by contract or otherwise, or (ii) ownership

of fifty percent (50%) or more of the outstanding shares, or (iii) beneficial ownership of such entity.

"You" (or "Your") shall mean an individual or Legal Entity exercising permissions granted by this License.

"Source" form shall mean the preferred form for making modifications, including but not limited to software source code, documentation source, and configuration files.

"Object" form shall mean any form resulting from mechanical transformation or translation of a Source form, including but not limited to compiled object code, generated documentation, and conversions to other media types.

"Work" shall mean the work of authorship, whether in Source or Object form, made available under the License, as indicated by a copyright notice that is included in or attached to the work (an example is provided in the Appendix below).

"Derivative Works" shall mean any work, whether in Source or Object form, that is based on (or derived from) the Work and for which the editorial revisions, annotations, elaborations, or other modifications represent, as a whole, an original work of authorship. For the purposes of this License, Derivative Works shall not include works that remain separable from, or merely link (or bind by name) to the interfaces of, the Work and Derivative Works thereof.

"Contribution" shall mean any work of authorship, including the original version of the Work and any modifications or additions to that Work or Derivative Works thereof, that is intentionally submitted to Licensor for inclusion in the Work by the copyright owner or by an individual or Legal Entity authorized to submit on behalf of the copyright owner. For the purposes of this definition, "submitted" means any form of electronic, verbal, or written communication sent to the Licensor or its representatives, including but not limited to communication on electronic mailing lists, source code control systems, and issue tracking systems that are managed by, or on behalf of, the Licensor for the purpose of discussing and improving the Work, but excluding communication that is conspicuously marked or otherwise designated in writing by the copyright owner as "Not a Contribution."

"Contributor" shall mean Licensor and any individual or Legal Entity on behalf of whom a Contribution has been received by Licensor and subsequently incorporated within the Work.

2. Grant of Copyright License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable copyright license to reproduce, prepare Derivative Works of, publicly display, publicly perform, sublicense, and distribute the Work and such Derivative Works in Source or Object form.
3. Grant of Patent License. Subject to the terms and conditions of this License, each Contributor hereby grants to You a perpetual, worldwide, non-exclusive, no-charge, royalty-free, irrevocable (except as stated in this section) patent license to make, have made, use, offer to sell, sell, import, and otherwise transfer the Work, where such license applies only to those patent claims licensable by such Contributor that are necessarily infringed by their Contribution(s) alone or by combination of their Contribution(s) with the Work to which such Contribution(s) was submitted. If You institute patent litigation against any entity (including a cross-claim or counterclaim in a lawsuit) alleging that the Work or a Contribution incorporated within the Work constitutes direct or contributory patent infringement, then any patent licenses granted to You under this License for that Work shall terminate as of the date such litigation is filed.
4. Redistribution. You may reproduce and distribute copies of the Work or Derivative Works thereof in any medium, with or without modifications, and in Source or Object form, provided that You meet the following conditions:
  - a. You must give any other recipients of the Work or Derivative Works a copy of this License; and
  - b. You must cause any modified files to carry prominent notices stating that You changed the files; and
  - c. You must retain, in the Source form of any Derivative Works that You distribute, all copyright, patent, trademark, and attribution notices from the Source form of the Work, excluding those notices that do not pertain to any part of the Derivative Works; and

- d. If the Work includes a "NOTICE" text file as part of its distribution, then any Derivative Works that You distribute must include a readable copy of the attribution notices contained within such NOTICE file, excluding those notices that do not pertain to any part of the Derivative Works, in at least one of the following places: within a NOTICE text file distributed as part of the Derivative Works; within the Source form or documentation, if provided along with the Derivative Works; or, within a display generated by the Derivative Works, if and wherever such third-party notices normally appear. The contents of the NOTICE file are for informational purposes only and do not modify the License. You may add Your own attribution notices within Derivative Works that You distribute, alongside or as an addendum to the NOTICE text from the Work, provided that such additional attribution notices cannot be construed as modifying the License.

You may add Your own copyright statement to Your modifications and may provide additional or different license terms and conditions for use, reproduction, or distribution of Your modifications, or for any such Derivative Works as a whole, provided Your use, reproduction, and distribution of the Work otherwise complies with the conditions stated in this License.

5. **Submission of Contributions.** Unless You explicitly state otherwise, any Contribution intentionally submitted for inclusion in the Work by You to the Licensor shall be under the terms and conditions of this License, without any additional terms or conditions. Notwithstanding the above, nothing herein shall supersede or modify the terms of any separate license agreement you may have executed with Licensor regarding such Contributions.
6. **Trademarks.** This License does not grant permission to use the trade names, trademarks, service marks, or product names of the Licensor, except as required for reasonable and customary use in describing the origin of the Work and reproducing the content of the NOTICE file.
7. **Disclaimer of Warranty.** Unless required by applicable law or agreed to in writing, Licensor provides the Work (and each Contributor provides its Contributions) on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied, including, without limitation, any warranties or conditions of TITLE, NON-INFRINGEMENT, MERCHANTABILITY, or FITNESS FOR A PARTICULAR PURPOSE. You are solely responsible for determining the appropriateness of using or redistributing the Work and assume any risks associated with Your exercise of permissions under this License.
8. **Limitation of Liability.** In no event and under no legal theory, whether in tort (including negligence), contract, or otherwise, unless required by applicable law (such as deliberate and grossly negligent acts) or agreed to in writing, shall any Contributor be liable to You for damages, including any direct, indirect, special, incidental, or consequential damages of any character arising as a result of this License or out of the use or inability to use the Work (including but not limited to damages for loss of goodwill, work stoppage, computer failure or malfunction, or any and all other commercial damages or losses), even if such Contributor has been advised of the possibility of such damages.
9. **Accepting Warranty or Additional Liability.** While redistributing the Work or Derivative Works thereof, You may choose to offer, and charge a fee for, acceptance of support, warranty, indemnity, or other liability obligations and/or rights consistent with this License. However, in accepting such obligations, You may act only on Your own behalf and on Your sole responsibility, not on behalf of any other Contributor, and only if You agree to indemnify, defend, and hold each Contributor harmless for any liability incurred by, or claims asserted against, such Contributor by reason of your accepting any such warranty or additional liability.

## END OF TERMS AND CONDITIONS

APPENDIX: How to apply the Apache License to your work.

To apply the Apache License to your work, attach the following boilerplate notice, with the fields enclosed by brackets "[ ]" replaced with your own identifying information. (Don't include the brackets!) The text should be enclosed in the appropriate comment syntax for the file format. We also recommend that a file or class name and description of purpose be included on the same "printed page" as the copyright notice for easier identification within third-party archives.

Copyright [yyyy] [name of copyright owner]



Licensed under the Apache License, Version 2.0 (the "License"); you may not use this file except in compliance with the License.

You may obtain a copy of the License at

<http://www.apache.org/licenses/LICENSE-2.0>

Unless required by applicable law or agreed to in writing, software distributed under the License is distributed on an "AS IS" BASIS, WITHOUT WARRANTIES OR CONDITIONS OF ANY KIND, either express or implied. See the License for the specific language governing permissions and limitations under the License.

## GNU Lesser General Public License

GNU LESSER GENERAL PUBLIC LICENSE

Version 2.1, February 1999

Copyright (C) 1991, 1999 Free Software Foundation, Inc.

59 Temple Place, Suite 330, Boston, MA 02111-1307 USA

Everyone is permitted to copy and distribute verbatim copies of this license document, but changing it is not allowed.

[This is the first released version of the Lesser GPL. It also counts as the successor of the GNU Library Public License, version 2, hence the version number 2.1.]

### Preamble

The licenses for most software are designed to take away your freedom to share and change it. By contrast, the GNU General Public Licenses are intended to guarantee your freedom to share and change free software--to make sure the software is free for all its users.

This license, the Lesser General Public License, applies to some specially designated software packages--typically libraries--of the Free Software Foundation and other authors who decide to use it. You can use it too, but we suggest you first think carefully about whether this license or the ordinary General Public License is the better strategy to use in any particular case, based on the explanations below.

When we speak of free software, we are referring to freedom of use, not price. Our General Public Licenses are designed to make sure that you have the freedom to distribute copies of free software (and charge for this service if you wish); that you receive source code or can get it if you want it; that you can change the software and use pieces of it in new free programs; and that you are informed that you can do these things.

To protect your rights, we need to make restrictions that forbid distributors to deny you these rights or to ask you to surrender these rights. These restrictions translate to certain responsibilities for you if you distribute copies of the library or if you modify it.

For example, if you distribute copies of the library, whether gratis or for a fee, you must give the recipients all the rights that we gave you. You must make sure that they, too, receive or can get the source code. If you link other code with the library, you must provide complete object files to the recipients, so that they can relink them with the library after making changes to the library and recompiling it. And you must show them these terms so they know their rights.

We protect your rights with a two-step method: (1) we copyright the library, and (2) we offer you this license, which gives you legal permission to copy, distribute and/or modify the library.

To protect each distributor, we want to make it very clear that there is no warranty for the free library. Also, if the library is modified by someone else and passed on, the recipients should know that what they have is not the original version, so that the original authors reputation will not be affected by problems that might be introduced by others.

Finally, software patents pose a constant threat to the existence of any free program. We wish to make sure that a company cannot effectively restrict the users of a free program by obtaining a restrictive license from a patent holder. Therefore, we insist that any patent license obtained for a version of the library must be consistent with the full freedom of use specified in this license.

Most GNU software, including some libraries, is covered by the ordinary GNU General Public License. This license, the GNU Lesser General Public License, applies to certain designated libraries, and is quite different from the ordinary

General Public License. We use this license for certain libraries in order to permit linking those libraries into non-free programs.

When a program is linked with a library, whether statically or using a shared library, the combination of the two is legally speaking a combined work, a derivative of the original library. The ordinary General Public License therefore permits such linking only if the entire combination fits its criteria of freedom. The Lesser General Public License permits more lax criteria for linking other code with the library.

We call this license the "Lesser" General Public License because it does Less to protect the users freedom than the ordinary General Public License. It also provides other free software developers Less of an advantage over competing non-free programs. These disadvantages are the reason we use the ordinary General Public License for many libraries. However, the Lesser license provides advantages in certain special circumstances.

For example, on rare occasions, there may be a special need to encourage the widest possible use of a certain library, so that it becomes a de-facto standard. To achieve this, non-free programs must be allowed to use the library. A more frequent case is that a free library does the same job as widely used non-free libraries. In this case, there is little to gain by limiting the free library to free software only, so we use the Lesser General Public License.

In other cases, permission to use a particular library in non-free programs enables a greater number of people to use a large body of free software. For example, permission to use the GNU C Library in non-free programs enables many more people to use the whole GNU operating system, as well as its variant, the GNU/Linux operating system.

Although the Lesser General Public License is Less protective of the users freedom, it does ensure that the user of a program that is linked with the Library has the freedom and the wherewithal to run that program using a modified version of the Library.

The precise terms and conditions for copying, distribution and modification follow. Pay close attention to the difference between a "work based on the library" and a "work that uses the library". The former contains code derived from the library, whereas the latter must be combined with the library in order to run.

## GNU LESSER GENERAL PUBLIC LICENSE

### TERMS AND CONDITIONS FOR COPYING, DISTRIBUTION AND MODIFICATION

0. This License Agreement applies to any software library or other program which contains a notice placed by the copyright holder or other authorized party saying it may be distributed under the terms of this Lesser General Public License (also called "this License"). Each licensee is addressed as "you".

A "library" means a collection of software functions and/or data prepared so as to be conveniently linked with application programs (which use some of those functions and data) to form executables.

The "Library", below, refers to any such software library or work which has been distributed under these terms. A "work based on the Library" means either the Library or any derivative work under copyright law: that is to say, a work containing the Library or a portion of it, either verbatim or with modifications and/or translated straightforwardly into another language. (Hereinafter, translation is included without limitation in the term "modification".)

"Source code" for a work means the preferred form of the work for making modifications to it. For a library, complete source code means all the source code for all modules it contains, plus any associated interface definition files, plus the scripts used to control compilation and installation of the library.

Activities other than copying, distribution and modification are not covered by this License; they are outside its scope. The act of running a program using the Library is not restricted, and output from such a program is covered only if its contents constitute a work based on the Library (independent of the use of the Library in a tool for writing it). Whether that is true depends on what the Library does and what the program that uses the Library does.

1. You may copy and distribute verbatim copies of the Library's complete source code as you receive it, in any medium, provided that you conspicuously and appropriately publish on each copy an appropriate copyright notice and disclaimer of warranty; keep intact all the notices that refer to this License and to the absence of any warranty; and distribute a copy of this License along with the Library.  
You may charge a fee for the physical act of transferring a copy, and you may at your option offer warranty protection in exchange for a fee.
2. You may modify your copy or copies of the Library or any portion of it, thus forming a work based on the Library, and copy and distribute such modifications or work under the terms of Section 1 above, provided that you also meet all of these conditions
  - a. The modified work must itself be a software library.

- b. You must cause the files modified to carry prominent notices stating that you changed the files and the date of any change.
- c. You must cause the whole of the work to be licensed at nocharge to all third parties under the terms of this License.
- d. If a facility in the modified Library refers to a function or a table of data to be supplied by an appropriate program that uses the facility, other than as an argument passed when the facility is invoked, then you must make a good faith effort to ensure that, in the event an application does not supply such function or table, the facility still operates, and performs whatever part of its purpose remains meaningful.

(For example, a function in a library to compute square roots has a purpose that is entirely well-defined independent of the application. Therefore, Subsection 2d requires that any application-supplied function or table used by this function must be optional: if the application does not supply it, the square root function must still compute square roots.)

These requirements apply to the modified work as a whole. If identifiable sections of that work are not derived from the Library, and can be reasonably considered independent and separate works in themselves, then this License, and its terms, do not apply to those sections when you distribute them as separate works. But when you distribute the same sections as part of a whole which is a work based on the Library, the distribution of the whole must be on the terms of this License, whose permissions for other licensees extend to the entire whole, and thus to each and every part regardless of who wrote it.

Thus, it is not the intent of this section to claim rights or contest your rights to work written entirely by you; rather, the intent is to exercise the right to control the distribution of derivative or collective works based on the Library. In addition, mere aggregation of another work not based on the Library with the Library (or with a work based on the Library) on a volume of a storage or distribution medium does not bring the other work under the scope of this License.

3. You may opt to apply the terms of the ordinary GNU General Public License instead of this License to a given copy of the Library. To do this, you must alter all the notices that refer to this License, so that they refer to the ordinary GNU General Public License, version 2, instead of to this License. (If a newer version than version 2 of the ordinary GNU General Public License has appeared, then you can specify that version instead if you wish.) Do not make any other change in these notices.

Once this change is made in a given copy, it is irreversible for that copy, so the ordinary GNU General Public License applies to all subsequent copies and derivative works made from that copy.

This option is useful when you wish to copy part of the code of the Library into a program that is not a library.

4. You may copy and distribute the Library (or a portion or derivative of it, under Section 2) in object code or executable form under the terms of Sections 1 and 2 above provided that you accompany it with the complete corresponding machine-readable source code, which must be distributed under the terms of Sections 1 and 2 above on a medium customarily used for software interchange.

If distribution of object code is made by offering access to copy from a designated place, then offering equivalent access to copy the source code from the same place satisfies the requirement to distribute the source code, even though third parties are not compelled to copy the source along with the object code.

5. A program that contains no derivative of any portion of the Library, but is designed to work with the Library by being compiled or linked with it, is called a "work that uses the Library". Such a work, in isolation, is not a derivative work of the Library, and therefore falls outside the scope of this License.

However, linking a "work that uses the Library" with the Library creates an executable that is a derivative of the Library (because it contains portions of the Library), rather than a "work that uses the library". The executable is therefore covered by this License. Section 6 states terms for distribution of such executables.

When a "work that uses the Library" uses material from a header file that is part of the Library, the object code for the work may be a derivative work of the Library even though the source code is not. Whether this is true is especially significant if the work can be linked without the Library, or if the work is itself a library. The threshold for this to be true is not precisely defined by law.

If such an object file uses only numerical parameters, data structure layouts and accessors, and small macros and small inline functions (ten lines or less in length), then the use of the object file is unrestricted, regardless of whether it is legally a derivative work. (Executables containing this object code plus portions of the Library will still fall under Section 6.)

Otherwise, if the work is a derivative of the Library, you may distribute the object code for the work under the terms of Section 6. Any executables containing that work also fall under Section 6, whether or not they are linked directly with the Library itself.

6. As an exception to the Sections above, you may also combine or link a "work that uses the Library" with the Library to produce a work containing portions of the Library, and distribute that work under terms of your choice, provided that the terms permit modification of the work for the customers own use and reverse engineering for debugging such modifications.

You must give prominent notice with each copy of the work that the Library is used in it and that the Library and its use are covered by this License. You must supply a copy of this License. If the work during execution displays copyright notices, you must include the copyright notice for the Library among them, as well as a reference directing the user to the copy of this License. Also, you must do one of these things:

- a. Accompany the work with the complete corresponding machine-readable source code for the Library including whatever changes were used in the work (which must be distributed under Sections 1 and 2 above); and, if the work is an executable linked with the Library, with the complete machine-readable "work that uses the Library", as object code and/or source code, so that the user can modify the Library and then relink to produce a modified executable containing the modified Library. (It is understood that the user who changes the contents of definitions files in the Library will not necessarily be able to recompile the application to use the modified definitions.)
- b. Use a suitable shared library mechanism for linking with the Library. A suitable mechanism is one that (1) uses at run time a copy of the library already present on the users computer system, rather than copying library functions into the executable, and (2) will operate properly with a modified version of the library, if the user installs one, as long as the modified version is interface-compatible with the version that the work was made with.
- c. Accompany the work with a written offer, valid for at least three years, to give the same user the materials specified in Subsection 6a, above, for a charge no more than the cost of performing this distribution.
- d. If distribution of the work is made by offering access to copy from a designated place, offer equivalent access to copy the above specified materials from the same place.
- e. Verify that the user has already received a copy of these materials or that you have already sent this user a copy.

For an executable, the required form of the "work that uses the Library" must include any data and utility programs needed for reproducing the executable from it. However, as a special exception, the materials to be distributed need not include anything that is normally distributed (in either source or binary form) with the major components (compiler, kernel, and so on) of the operating system on which the executable runs, unless that component itself accompanies the executable.

It may happen that this requirement contradicts the license restrictions of other proprietary libraries that do not normally accompany the operating system. Such a contradiction means you cannot use both them and the Library together in an executable that you distribute.

7. You may place library facilities that are a work based on the Library side-by-side in a single library together with other library facilities not covered by this License, and distribute such a combined library, provided that the separate distribution of the work based on the Library and of the other library facilities is otherwise permitted, and provided that you do these two things:
  - a. Accompany the combined library with a copy of the same work based on the Library, uncombined with any other library facilities. This must be distributed under the terms of the Sections above.
  - b. Give prominent notice with the combined library of the fact that part of it is a work based on the Library, and explaining where to find the accompanying uncombined form of the same work.
8. You may not copy, modify, sublicense, link with, or distribute the Library except as expressly provided under this License. Any attempt otherwise to copy, modify, sublicense, link with, or distribute the Library is void, and will automatically terminate your rights under this License. However, parties who have received copies, or rights, from you under this License will not have their licenses terminated so long as such parties remain in full compliance.
9. You are not required to accept this License, since you have not signed it. However, nothing else grants you permission to modify or distribute the Library or its derivative works. These actions are prohibited by law if you do not accept this License. Therefore, by modifying or distributing the Library (or any work based on the Library), you indicate your acceptance of this License to do so, and all its terms and conditions for copying, distributing or modifying the Library or works based on it.
10. Each time you redistribute the Library (or any work based on the Library), the recipient automatically receives a license from the original licensor to copy, distribute, link with or modify the Library subject to these terms and conditions. You may not impose any further restrictions on the recipients exercise of the rights granted herein. You are not responsible for enforcing compliance by third parties with this License.
11. If, as a consequence of a court judgment or allegation of patent infringement or for any other reason (not limited to patent issues), conditions are imposed on you (whether by court order, agreement or otherwise) that contradict the conditions of this License, they do not excuse you from the conditions of this License. If you cannot distribute so as to satisfy simultaneously your obligations under this License and any other pertinent obligations, then as a consequence you may not distribute the Library at all. For example, if a patent license would not permit royalty-free redistribution of the Library by all those who receive copies directly or indirectly through you, then the only way you could satisfy both it and this License would be to refrain entirely from distribution of the Library. If any portion of this section is held invalid or unenforceable under any particular circumstance, the balance of the

section is intended to apply, and the section as a whole is intended to apply in other circumstances. It is not the purpose of this section to induce you to infringe any patents or other property right claims or to contest validity of any such claims; this section has the sole purpose of protecting the integrity of the free software distribution system which is implemented by public license practices. Many people have made generous contributions to the wide range of software distributed through that system in reliance on consistent application of that system; it is up to the author/donor to decide if he or she is willing to distribute software through any other system and a licensee cannot impose that choice.

This section is intended to make thoroughly clear what is believed to be a consequence of the rest of this License.

12. If the distribution and/or use of the Library is restricted in certain countries either by patents or by copyrighted interfaces, the original copyright holder who places the Library under this License may add an explicit geographical distribution limitation excluding those countries, so that distribution is permitted only in or among countries not thus excluded. In such case, this License incorporates the limitation as if written in the body of this License.
13. The Free Software Foundation may publish revised and/or new versions of the Lesser General Public License from time to time. Such new versions will be similar in spirit to the present version, but may differ in detail to address new problems or concerns.  
Each version is given a distinguishing version number. If the Library specifies a version number of this License which applies to it and "any later version", you have the option of following the terms and conditions either of that version or of any later version published by the Free Software Foundation. If the Library does not specify a license version number, you may choose any version ever published by the Free Software Foundation.
14. If you wish to incorporate parts of the Library into other free programs whose distribution conditions are incompatible with these, write to the author to ask for permission. For software which is copyrighted by the Free Software Foundation, write to the Free Software Foundation; we sometimes make exceptions for this. Our decision will be guided by the two goals of preserving the free status of all derivatives of our free software and of promoting the sharing and reuse of software generally.

#### NO WARRANTY

15. BECAUSE THE LIBRARY IS LICENSED FREE OF CHARGE, THERE IS NO WARRANTY FOR THE LIBRARY, TO THE EXTENT PERMITTED BY APPLICABLE LAW. EXCEPT WHEN OTHERWISE STATED IN WRITING THE COPYRIGHT HOLDERS AND/OR OTHER PARTIES PROVIDE THE LIBRARY "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE. THE ENTIRE RISK AS TO THE QUALITY AND PERFORMANCE OF THE LIBRARY IS WITH YOU. SHOULD THE LIBRARY PROVE DEFECTIVE, YOU ASSUME THE COST OF ALL NECESSARY SERVICING, REPAIR OR CORRECTION.
16. IN NO EVENT UNLESS REQUIRED BY APPLICABLE LAW OR AGREED TO IN WRITING WILL ANY COPYRIGHT HOLDER, OR ANY OTHER PARTY WHO MAY MODIFY AND/OR REDISTRIBUTE THE LIBRARY AS PERMITTED ABOVE, BE LIABLE TO YOU FOR DAMAGES, INCLUDING ANY GENERAL, SPECIAL, INCIDENTAL OR CONSEQUENTIAL DAMAGES ARISING OUT OF THE USE OR INABILITY TO USE THE LIBRARY (INCLUDING BUT NOT LIMITED TO LOSS OF DATA OR DATA BEING RENDERED INACCURATE OR LOSSES SUSTAINED BY YOU OR THIRD PARTIES OR A FAILURE OF THE LIBRARY TO OPERATE WITH ANY OTHER SOFTWARE), EVEN IF SUCH HOLDER OR OTHER PARTY HAS BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

#### END OF TERMS AND CONDITIONS

## Product Compliance

This section describes the hardware components, requirements, and specifications, plus agency and RFC (Request for Comments) compliance for the Infoblox appliance. Topics in this section include:

- [RFC Compliance](#)
- [Agency Compliance](#)
- [DC](#)
- [AC](#)

## Power Safety Information

The main external power connector for the Infoblox appliance is located on the back of the system. Ensure power to the system is off before connecting the power cord into the power connector. Please read the following power safety statements for your AC- or DC-powered appliance:

## RFC Compliance

The NIOS appliance is compliant with the following:

- Qualys and Nessus security requirements
- Joint Interoperability Test Command (JITC) certification for Internet Protocol version 6 capability
- RFCs (Request for Comments)

The list of RFCs are:

- [DNS RFC Compliance](#)
- [DHCP RFC Compliance](#)
- [DHCPv6 RFC Compliance](#)
- [IDN \(Internationalized Domain Names\) RFC Compliance](#)

## DNS RFC Compliance

The NIOS appliance complies with the following DNS RFCs:

*RFCs for DNS*

RFC Number	RFC Title
805	Computer Mail Meeting Notes
811	Hostnames Server
819	The Domain Naming Convention for Internet User Applications
881	The Domain Names Plan and Schedule
882	Domain Names: Concepts and Facilities
883	Domain Names: Implementation Specification
897	Domain Name System Implementation Schedule
920	Domain Requirements
921	Domain Name System Implementation Schedule – Revised
973	Domain System Changes and Observations

<b>RFC Number</b>	<b>RFC Title</b>
974	Mail Routing and the Domain System
1032	Domain Administrators Guide
1033	Domain Administrators Operations Guide
1034	Domain Names – Concepts and Facilities
1035	Domain Names – Implementation and Specification
1101	DNS Encoding of Network Names and Other Types
1122	Requirements for Internet Hosts – Communication Layers
1123	Requirements for Internet Hosts – Application and Support
1178	Choosing a Name for Your Computer
1348	DNS NSAP RRs
1386	The US Domain
1464	Using the Domain Name System to Store Arbitrary String Attributes
1535	A Security Problem and Proposed Correction with Widely Deployed DNS Software
1536	Common DNS Implementation Errors and Suggested Fixes
1537	Common DNS Data File Configuration Errors
1591	Domain Name System Structure and Delegation
1611	DNS Server MIB Extensions
1612	DNS Resolver MIB Extensions
1637	DNS NSAP Resource Records
1664	Using the Internet DNS to Distribute RFC
1327	Mail Address Mapping Tables
1713	Tools for DNS debugging

<b>RFC Number</b>	<b>RFC Title</b>
1794	DNS Support for Load Balancing
1811	U.S. Government Internet Domain Names
1816	U.S. Government Internet Domain Names
1912	Common DNS Operational and Configuration Errors
1956	Registration in the MIL Domain
1982	Serial Number Arithmetic
1995	Incremental Zone Transfer in DNS
1996	A Mechanism for Prompt Notification of Zone Changes
2010	Operational Criteria for Root Name Servers
2052	A DNS RR for specifying the location of services (DNS SRV)
2053	The AM (Armenia) Domain
2136	Dynamic Updates in the Domain Name System (DNS UPDATE)
2142	Mailbox Names for Common Services, Roles and Functions
2146	U.S. Government Internet Domain Names
2168	Resolution of Uniform Resource Identifiers using the Domain Name System
2181	Clarifications to the DNS Specification
2182	Selection and Operation of Secondary DNS Servers
2219	Use of DNS Aliases for Network Services
2240	A Legal Basis for Domain Name Allocation
2308	Negative Caching of DNS Queries (DNS NCACHE)
2317	Classless IN-ADDR.ARPA Delegation
2352	A Convention for Using Legal Names as Domain Names



RFC Number	RFC Title
2537	RSA/MD5 KEYS and SIGs in the Domain Name System (DNS)
2606	Reserved Top Level DNS Names
2782	A DNS RR for Specifying the Location of Services (DNS SRV)
2845	Secret Key Transaction Authentication for DNS (TSIG)
2915	The Naming Authority Pointer (NAPTR) DNS Resource Record
2930	Secret Key Establishment for DNS (TKEY RR)
3596	DNS Extensions to Support IP Version 6
3645	Generic Security Service Algorithm for Secret Key Transaction Authentication for DNS (GSS-TSIG)
3768	Virtual Router Redundancy Protocol (VRRP)
4033	DNS Security Introduction and Requirements
4034	Resource Records for the DNS Security Extensions
4035	Protocol Modifications for the DNS Security Extensions
4641	DNSSEC Operational Practices
4956	DNS Security (DNSSEC) Opt-In
4986	Requirements Related to DNS Security (DNSSEC) Trust Anchor Rollover
5155	DNSSEC Hashed Authenticated Denial of Existence
5702	Use of SHA-2 Algorithms with RSA in DNSKEY and RRSIG Resource Records for DNSSEC
5936	DNS Zone Transfer Protocol (AXFR)
6147	DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers
6698	The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA
6844	DNS Certification Authority Authorization (CAA) Resource Record
6891	Extension Mechanisms for DNS (EDNS0)

RFC Number	RFC Title
7646	Definition and Use of DNSSEC Negative Trust Anchors
7671	The DNS-Based Authentication of Named Entities (DANE) Protocol
7766	DNS Transport over TCP
7871	Client Subnet in DNS Queries

## DHCP RFC Compliance

The appliance complies with the following DHCP RFCs:

*RFCs for DHCP*

RFC Number	RFC Title
1531	Dynamic Host Configuration Protocol 1534 Interoperation Between DHCP and BOOTP
1542	Clarifications and Extensions for the Bootstrap Protocol
2131	Dynamic Host Configuration Protocol
2132	DHCP Options and BOOTP Vendor Extensions
3046	DHCP Relay Agent Information Option
3315	Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
3925	Vendor-Identifying Vendor Options for Dynamic Host Configuration Protocol version 4 (DHCPv4)
4388	Dynamic Host Configuration Protocol (DHCP) Leasequery

## DHCPv6 RFC Compliance

The appliance complies with the following DHCPv6 RFCs:

*RFCs for DHCPv6*

RFC Number	RFC Title
4075	Simple Network Time Protocol (SNTP) Configuration Option for DHCPv6
3898	Network Information Service (NIS) Configuration Options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
3736	Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6

RFC Number	RFC Title
3646	DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
3633	IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6
3319	Dynamic Host Configuration Protocol (DHCPv6) Options for Session Initiation Protocol (SIP) Servers

## IDN (Internationalized Domain Names) RFC Compliance

The appliance complies with the following IDN RFCs:

### *RFCs for IDN*

RFC Number	RFC Title
3492	Punycode: A Bootstring encoding of Unicode for Internationalized Domain Names in Applications (IDNA)
5890	Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework
5891	Internationalized Domain Names for Applications (IDNA): Background, Explanation, and Rationale
5892	The Unicode code points and IDNA
5893	Right-to-left scripts for IDNA
5894	Internationalized Domain Names in Applications (IDNA): Protocol
5895	Mapping Characters in IDNA2008
6452	The Unicode Code Points and Internationalized Domain Names for Applications (IDNA) - Unicode 6.0

## Agency Compliance

The Infoblox appliance is compliant with these EMI and safety agency regulations:

### *Agency Regulation Compliance*

Standard	Agency	Marks
FCC Part 15	FCC	FCC
EN55022, EN55024, EN61000-3-2, EN61000-3-3	TUV	CE
UL60950/CSA60950	UL	cULus
EN60950	TUV	GS

Standard	Agency	Marks
CB Scheme	IECEE	Report and Certificate IEC 60950-1:2001
VCCI-A	VCCI	VCCI
AS/NZS 3548	ACMA	C-Tick

## FCC

The FCC label on the back of the system indicates this network appliance is compliant with limits for a Class A digital device in accordance with Part 15 of the FCC Rules. These limits are designed to provide reasonable protection against harmful interference when this equipment is operated in a commercial environment. Operation is subject to the following two conditions:

- This device might not cause harmful interference.
- This device must accept any interference received, including interference that may cause undesired operation.

This device generates, uses, and can radiate radio frequency energy if not installed and used in accordance with the instructions in this manual. Operating this equipment in a residential area is likely to cause harmful interference, and the customer will be required to rectify the interference at his or her own expense. This product requires the use of external shielded cables to maintain compliance pursuant to Part 15 of the FCC Rules.

## Canadian Compliance

### English

This Class A digital apparatus complies with Canadian ICES-003.

### French

Cet appareil numérique de la classe A est conforme à la norme NMB-003 du Canada.

## VCCI

The Infoblox appliance complies with this VCCI regulation (compliance statement follow by its translation):

この装置は、情報処理装置等電波障害自主規制協議会（VCCI）の基準に基づくクラスA情報技術装置です。この装置を家庭環境で使用すると電波妨害を引き起こすことがあります。この場合には使用者が適切な対策を講ずるよう要求されることがあります。

This is a Class A product based on the Technical Requirements of the Voluntary Control Council for Interference Technology (VCCI). In a domestic environment this product may cause radio interference, in which case the user may be required to take corrective action.



### Caution

Lithium battery included with this board. Do not puncture, mutilate, or dispose of battery in fire. Danger of explosion if battery is incorrectly replaced. Replace only with the same or equivalent type recommended by manufacturer. Dispose of used battery according to manufacturer instructions and in accordance with your local regulations.

## DC

### English

- ⚠ When stranded wiring is required, use approved wiring terminations, such as closed-loop or spade-type with upturned lugs. These terminations should be the appropriate size for the wires and should clamp both the insulation and conductor.

## AC

### English

- ⚠ **Warning**  
This product relies on the building's installation for short-circuit (overcurrent) protection. Ensure that a fuse or circuit breaker no larger than 120VAC, 15A U.S. (240VAC, 10A international) is used on the phase conductors (all current-carrying conductors)

### French

- ⚠ **Warning**  
Pour ce qui est de la protection contre les courts-circuits (surtension), ce produit dépend de l'installation électrique du local. Vérifier qu'un fusible ou qu'un disjoncteur de 120V alt., 15A U.S. maximum (240V alt., 10A international) est utilisé sur les conducteurs de phase (conducteurs de charge).

### German

- ⚠ **Warning**  
Dieses Produkt ist darauf angewiesen, daß im Gebäude ein Kurzschluß - bzw. Überstromschutz installiert ist. Stellen Sie sicher, daß eine Sicherung oder ein Unterbrecher von nicht mehr als 240V Wechselstrom, 10A (bzw. in den USA 120V Wechselstrom, 15A) an den Phasenleitern (allen stromführenden Leitern) verwendet wird.

## vNIOs Appliances

vNIOs appliances support most of the features of the Infoblox NIOs software, with some limitations. The subtopics in this section describe these limitations.

- [vNIOs for Appliances Limitations](#)



## Notes

- VMware Tools are automatically installed for each vNIOs appliance. Infoblox supports the control functions in Hyper-V Manager and VMware Tools. For example, through the vSphere client, you can shut down the virtual appliance.
- The VMXNET virtual network adapter for vNIOs is not supported from NIOS 8.4.x onwards.

Infoblox NIOS virtual appliances support any hardware that provides the required Hypervisor version, memory, CPU, and disk resources. To maintain high performance on your NIOS virtual appliances and to avoid not having enough resources to service all the NIOS virtual appliances, do not oversubscribe physical resources on the virtualization host. Required memory, CPU, and disk resources must be adequately allocated for each virtual appliance that is running on the virtualization host. For information about the required specification for each NIOS virtual appliance model, see the following table:

*Supported vNIOs Appliance Models and Specifications*

NIOS Virtual Appliance s	Overall Disk (GB)	# of CPU Cores	Memory Allocation (GB)	Recommended CPU Per Core Clock Rate	NIOS for VMware	NIOS for MS Hyper-V*	NIOS for KVM	NIOS for AWS, Azure and GCP	NIOS for Nutanix AHV	NIOS for Red Hat Open Shift	Supported as Grid Master and Grid Master Candidate
IB-V815**	250	2	16	1100 MHz	✓	✓	✓ <sup>1</sup>	X	✓	X	Yes
IB-V825**	250	2	16	1600 MHz	✓	✓	✓ <sup>1</sup>	✓	✓	X	Yes
IB-V1415**	250	4	32	1200 MHz	✓	✓	✓ <sup>1</sup>	X	✓	X	Yes
IB-V1425**	250	4	32	1800 MHz	✓	✓	✓ <sup>1</sup>	✓	✓	X	Yes
IB-V2215**	250	8	64	2100 MHz	✓	✓	✓ <sup>1</sup>	X	✓	X	Yes
IB-V2225**	250	8	64	2100 MHz	✓	✓	✓ <sup>1</sup>	✓	✓	✓	Yes
IB-V4015**	250	14	128	2400 MHz	✓	✓	✓ <sup>1</sup>	✓ <sup>2</sup>	X	✓	Yes
IB-V4025**	250	16	122	2400 MHz	✓	✓	✓ <sup>1</sup>	✓ <sup>2</sup>	X	X	Yes

NIOS Reporting Virtual Appliances	Overall Disk (GB)	# of CPU Cores	Memory Allocation (GB)	Recommended CPU Per Core Clock Rate	NIOS for VMware	NIOS for MS Hyper-V*	NIOS for KVM	NIOS for AWS, Azure	NIOS for GCP	NIOS for Nutanix AHV	Supported as Grid Master and Grid Master Candidate
IB-V805**	250	2	32	2700 MHz	✓	✓	✓ <sub>1</sub>	X	X	X	No
IB-V1405**	250	4	32	3600 MHz	✓	✓	✓ <sub>1</sub>	X	X	X	No
IB-V2205**	250	8	64	2100 MHz	✓	✓	✓ <sub>1</sub>	X	X	X	No
IB-V4005	250 (+ 1500 GB reporting storage)	14	128	2400 MHz	✓	X	X	X	X	X	No
IB-V5005	User defined reporting storage	User defined	User defined	N / A	✓	✓	✓	✓ <sub>3</sub>	X	✓	No

Network Insight Virtual Appliances	Overall Disk (GB)	# of CPU Cores	Memory Allocation (GB)	Recommended CPU Per Core Clock Rate	NIOS for AWS, Azure, and VMware***	NIOS for MS Hyper-V*	NIOS for KVM	NIOS for GCP	NIOS for Nutanix AHV	Supported as Grid Master and Grid Master Candidate
ND-V805**	500	2	32	2700 MHz	✓	✓	✓ <sub>1</sub>	X	X	No
ND-V1405**	500	8	32	3600 MHz	✓	✓	✓ <sub>1</sub>	X	✓	No
ND-V2205**	500	16	64	2100 MHz	✓	✓	✓	X	X	No
ND-V4005**	500	28	128	2400 MHz	✓	✓	✓	X	X	No

Cloud Platform Appliances	Overall Disk (GB)	# of CPU Cores	Memory Allocation (GB)	Recommended CPU Per Core Clock Rate	NIOS for VMware	NIOS for MS Hyper-V*	NIOS for KVM	NIOS for AWS, Azure, GCP	NIOS for Nutanix AHV	NIOS for Oracle Cloud Infrastructure	Supported as Grid Master and Grid Master Candidate
CP-V805	250	2	16	2000 MHz	✓	✓	✓	✓	✓	X	No
CP-V1405	250	4	32	6000 MHz	✓	✓	✓	✓	✓	X	No
CP-V2205	250	8	64	12000 MHz	✓	✓	✓	✓	✓	✓	No

\* When running NIOS in MS Hyper-V with dynamic memory allocation enabled, your system might experience high memory usage. To avoid this issue, Infoblox recommends that you disable dynamic memory allocation.

\* For optimal performance, vNIOS for Hyper-V is not recommended as a Grid Master or Grid Master Candidate.

\* Specifications of vNIOS for Microsoft Azure Stack Hub are different from the other vNIOS for Microsoft Azure flavors. For the exact specifications, see [Infoblox Installation Guide vNIOS for Microsoft Azure](#).

\*\* To achieve best performance on your virtual appliances, follow the recommended specifications and allocate your resources within the limits of the licenses being installed on the appliances.

\*\* Specifications of Network Insight appliances for vNIOS for AWS and vNIOS for Azure are different from other vNIOS products. For detailed specifications, see the [Infoblox Installation Guide for vNIOS for AWS](#) and [Infoblox Installation Guide vNIOS for Microsoft Azure](#) respectively.

<sup>1</sup> NIOS for KVM is supported in the following environments: Red Hat OpenStack and Ubuntu. Note that only IB-V1405 as a reporting server has been qualified for Red Hat OpenStack.

<sup>2</sup> vNIOS for AWS is supported on the IB-V4025 appliance running NIOS 8.5.2 or later, and IB-V4015 appliance running NIOS 8.6.2 or later. vNIOS for Azure and vNIOS for GCP instances are supported on IB-V4015 and IB-V4025 appliances running NIOS 8.6.2 or later.

<sup>3</sup> vNIOS for AWS and vNIOS for Azure instances are supported only on the IB-V5005 appliance running NIOS 8.6.2 or later. For detailed specifications, see the [Infoblox Installation Guide for vNIOS for AWS](#) and [Infoblox Installation Guide vNIOS for Microsoft Azure](#).

For information about the limitations on each vNIOS appliance, see [vNIOS for VMware Limitations](#).

NIOS 8.6.1 is not supported on the following appliances:

IB-250, IB-250-A, IB-500, IB-550, IB-550-A, IB-1000, IB-1050, IB-1050-A, IB-1550, IB-1550-A, IB-1552, IB-1552-A, IB-1852-A, IB-2000, IB-2000-A, IB-VM-250, IB-VM-550, IB-VM-1050, IB-VM-1550, IB-VM-1850, IB-VM-2000, PT-1400, PT-2200, PT-4000, PT-4000-10GE, ND-800, ND-1400, ND-2200, TE-100, TE-810, TE-820, TE-1410, TE-1420, TE-2210, TE-2220, TR-800, TR-1400, TR-2200, IB-4010, IB-4020, TR-4000, IB-4030, ND-V800, ND-V1400, ND-V2200, TE-V810, TE-V820, TE-V1410, TE-V1420, TE-V2210, TE-V2220, TR-V2200, IB-V4010, IB-V4020, TE-V800, TE-V1400, TE-V2200, CP-V800, CP-V1400, CP-V2200, and TrinziC Reporting TR-2000 and TR-2000-A series appliances.

Because these appliances are not supported, Infoblox recommends that you do not upgrade to NIOS 8.6.1 on these appliances.



### Note

- TE appliances are also known as the IB appliances.



- You can freely assign resources to these virtual appliances to suit your business needs, as long as the resources are within the limits of the licenses being installed on the appliances.
- For the latest and most complete list of supported platforms, see the 8.6.1 Release Notes.

## vNIOs for Appliances Limitations

This topic lists the limitations on each vNIOs appliance:

- [vNIOs for VMWare Limitations](#)
- [vNIOs for Hyper-V Limitations](#)

## vNIOs for VMWare Limitations

The Infoblox vNIOs for VMware can also run on Cisco SRE-V (Services Ready Engine Virtualization), which is part of the Cisco UCS (Unified Computing System) Express. For more information about vNIOs for VMware, refer to the *Infoblox Installation Guide for vNIOs Software on VMware*.

vNIOs for VMware appliances support most of the features of the Infoblox NIOs appliances, with the following limitations:

- You must have a vNIOs license installed on the appliance before you can access the Infoblox GUI.
- vNIOs appliances do not support the following features:
  - Configuration of port settings for MGMT, LAN, LAN2, and HA ports
  - The bloxTools environment
- The IB-BOB virtual appliance is supported on Cisco SRE-V and can function as a Grid member only. It does not support configuration as an independent appliance, an HA pair, a Grid Master, or a Grid Master candidate. It also does not support access to the Infoblox GUI.
- The IB-VM-250 virtual appliance supports all the services provided by vNIOs virtual appliances, but it is not recommended as a Grid Master or Grid Master candidate.
- If you are using NIOs versions earlier than 8.6.x, if you configure an HA pair, both nodes in the HA pair must be NIOs virtual instances. You cannot configure a physical NIOs appliance and a NIOs virtual instance in an HA pair in versions earlier than 8.6.x.
- vNIOs appliances run on virtual hardware. They do not have sensors to monitor the physical CPU temperature, fan speed, and system temperature.
- Changing the vNIOs appliance settings through the VMware vSphere or vCenter console may violate the terms of the vNIOs licensing and support models. The vNIOs appliance may not join the Grid or function properly.

## vNIOs for Hyper-V Limitations

vNIOs for Microsoft Windows 2008 R2 and 2012 R2 server appliances support most of the features of the Infoblox NIOs appliances, with the following limitations:

- You must have a vNIOs license installed on the appliance before you can access the Infoblox GUI.
- vNIOs appliances do not support the following features:
  - Configuration of port settings for MGMT, LAN, LAN2, and HA ports
  - The bloxTools environment
- All the IB-VM appliance models support all the services provided by vNIOs virtual appliances, but Grid Master or Grid Master Candidate is not supported.
- The Captive Portal is supported only on IB-VM-1410 virtual appliances.
- vNIOs appliances run on virtual hardware. They do not have sensors to monitor the physical CPU temperature, fan speed, and system temperature.
- Changing the vNIOs appliance settings through the Hyper-V Manager or Virtual Machine Manager Administrator console may violate the terms of the vNIOs licensing and support models. The vNIOs appliance may not join the Grid or function properly.

# Guidance Documentation

Common Criteria provides an independent and objective evaluation of the security of Information Technology (IT) products. It gives assurance that the product satisfies a set of internationally recognized security standards.

This section provides additional guidance on the secure installation of the Target of Evaluation (TOE) for Common Criteria Evaluation Assurance Level (EAL) 2. To ensure that your appliance is Common Criteria compliant, make sure that your hardware and software settings match the evaluated configuration that was certified for Common Criteria.

This section should be used as the guiding document for installation of the TOE in the Common Criteria evaluated configuration.



## Note

Standalone NIOS is the Common Criteria compliant configuration for the TOE, which does not support Grid configuration.

This section contains the following topics:

- [Secure Syslog Transport](#)
- [Audit Log](#)
- [Backing Up and Restoring the Database](#)
- [About DNS](#)
- [WebUI Settings](#)
- [Licenses and Services](#)
- [Enabling / Disabling Common Criteria Mode](#)
- [Administration](#)
- [Security Guidelines](#)
- [Syslog](#)

## Secure Syslog Transport

The TOE provides the secure syslog transport feature using the TLS protocol. This allows secure transmission of messages between the syslog client, i.e. your NIOS appliance, and an external syslog server. You must use this feature to be Common Criteria compliant.

To ensure secure syslog transport, you add a trusted CA certificate of the server to your NIOS appliance. The certificate is then used to establish a secure connection to the server before transmitting data. For more information, see [Specifying Syslog Servers](#).

## Audit Log

The audit log contains a record of all TOE administrative activities. The stored audit records in the audit trail are protected from unauthorized modifications and deletion. For more information about the audit log, see [Using the Audit Log](#). Following are the events that are logged and examples of their corresponding audit log messages:

### Identification and Authentication

**Event:** Invalid password when logging in to the WebUI.

**Message:** "2011-10-19 14:02:32.750Z [admin]: Login\_Denied - - to=Serial\040Console apparently\_via=Directerror=invalid\040login\040or\040password"

**Event:** Number of attempts exceeds the limit when logging in to the WebUI.  
**Message:** "2011-10-19 14:05:23.217Z [admin]: Login\_Denied - - to=Serial\040Console apparently\_via=Directerror=failed\040logins\040exceed\040limit"

**Event:** Invalid password when logging in to the CLI.  
**Message:** "2011-10-19 14:02:32.750Z [admin]: Login\_Denied - - to=Serial\040Console apparently\_via=Directerror=invalid\040login\040or\040password"

**Event:** Number of attempts exceeds the limit when logging in to the CLI.  
**Message:** "2011-10-19 14:05:23.217Z [admin]: Login\_Denied - - to=Serial\040Console apparently\_via=Directerror=failed\040logins\040exceed\040limit"

**Event:** Enable Common Criteria mode:  
**Message:** 2011-10-19 19:48:37.299Z [admin]: Login\_Allowed - - to=Serial\040Console apparently\_via=Directauth=Local group=.admin-group

**Message:** 2011-10-19 19:48:48.705Z [admin]: Called - set\_cc\_mode: Args cc\_mode\_enabled="true"

**Event:** Disable Common Criteria mode:

**Message:** 2011-10-19 19:48:37.299Z [admin]: Login\_Allowed - - to=Serial\040Console apparently\_via=Directauth=Local group=.admin-group

**Message:** 2011-10-19 19:48:48.705Z [admin]: Called - set\_cc\_mode: Args cc\_mode\_enabled="false"

**Event:** Login successful

**Message:** 2011-10-19 19:48:48.706Z [USER\040admin]: rebooted the system  
2011-11-01 17:09:21.696Z [admin]: Login\_Allowed - - to=Serial\040Console apparently\_via=Direct auth=Localgroup=.admin-group

**Event:** First login

**Message:** 2011-10-19 12:43:47.375Z [user]: First\_Login - - to=AdminConnector ip=127.0.0.1 auth=LOCALgroup=admin-group apparently\_via=GUI first login

**Event:** Password expired

**Message:** 2011-10-20 13:17:29.257Z [user]: Password\_Expired - - to=AdminConnector ip=127.0.0.1 auth=LOCALgroup=admin-group apparently\_via=GUI

**Event:** Password was successfully reset.

**Message:** 2011-10-19 12:44:45.962Z [user]: Password\_Reset - - to=AdminConnector auth=LOCALgroup=admin-group apparently\_via=GUI

**Event:** New password did not conform to the rule.

**Message:** 2011-10-19 13:07:33.343Z [user]: Password\_Reset\_Error - - to=AdminConnector auth=LOCALgroup=admin-group apparently\_via=GUI

## Quotas

**Event:** Upload file limit reached.

**Message:** user manojk-vm httpd[]: err User {0} tried to upload the file. File {1} with size 272629904 kBytes is greater than maximum size allowed. Maximum size is 102400 kBytes.

## LDAP

**Event:** Establishment of session

**Message:** 2011-10-27T07:50:59-04:00 user epbyminw0065t2 python[]: notice Connection established:success

**Event:** Failure to establish a session

**Message:** 2011-10-27T07:50:38-04:00 user epbyminw0065t2 python[]: err 10.6.11.249: AD user authentication timed out

**Message:** 2011-10-27T07:51:02-04:00 user epbyminw0065t2 python[]: err Connection timed out

**Event:** Crypto Failure (Type and name of crypto algorithm that failed cannot be logged, since openldap uses SSL/TLS protocol functions from OpenSSL and did not use crypto functions directly.)

**Message:** 2011-10-27T07:51:00-04:00 user epbyminw0065t2 python[]: err SSL handshake failed.

**Message:** 2011-10-27T07:51:02-04:00 user epbyminw0065t2 python[]: err SSL handshake failed. Cannot verify server certificate.

## GSS-TSIG

**Event:** Invalid size specified for algorithm HMAC-SHA256

**Message:** 2011-10-19T17:57:12-04:00 user EPBYMINW2856 httpd[]: err TSIG key generation failure: Size 512 can not be used with algorithm HMAC-SHA256

**Event:** Invalid algorithm specified in Common Criteria mode

**Message:** 2011-10-19T18:12:22-04:00 user EPBYMINW2856 httpd[]: err TSIG key (keylen = 256, alname = HMAC-MD5) generation error : Only HMAC-SHA256 available in CC mode.

**Event:** Algorithm restriction

**Message:** Only AES128\_CTS\_HMAC\_SHA1\_96 or AES256\_CTS\_HMAC\_SHA1\_96 algorithms are allowed in CC mode. Current algorithm is DES\_CBC\_CRC.

## TSIG CSV Import/Export

**Event:** Import error (TSIG algorithm is not allowed in Common Criteria mode)

**Message:** [2011/10/20 09:38:42.496] (24473 /usr/bin/python)/infoblox/common/lib/python/infoblox/one/

csv\_import\_function.py:601 write\_to\_error\_file(): Import

Error:authzone,[zone.com](http://zone.com),FORWARD,,,,,,,,False,False,False,,1.2.3.4/1.2.3.4/False/False/True/ext\_sec\_key/ut29ROLaJwty6a%2Fhsgg0wA==,infoblox.localdomain,False,,,,,,,,,,,,,2,,default,Authoritative-Line 2: Insertion aborted due to IBDataError?: IB.Data:TSIG algorithm used for TSIG key name 'ext\_sec\_key' is not allowed in CC mode.

## “set” Commands

**Message:** 2011-10-19 13:14:04.030Z [admin]: Called - set\_snmptrap: Args variable="sysName.0", address="10.120.20.31"

**Message:** 2011-10-19 13:16:16.545Z [admin]: Called - set\_scheduled: Args task\_restarts="0 from 60"

**Message:** 2011-10-19 13:17:19.391Z [admin]: Called - set\_mld\_version\_1: MLD version set to 1

**Message:** 2011-10-19 13:18:28.171Z [admin]: Called - set\_support\_access: Args support\_access="true from false"

**Message:** 2011-10-19 13:19:46.669Z [admin]: Called - set\_session\_timeout: Args session\_timeout="650 from 600"

**Message:** 2011-10-19 13:23:11.596Z [admin]: Called - set\_phonehome: Args phonehome\_disabled="true from false"

**Message:** 2011-10-19 13:24:02.372Z [admin]: Called - set\_remote\_console: Args remote\_console="true from false"

**Message:** 2011-10-19 13:25:31.704Z [admin]: Called - set\_security: Args

address="10.120.20.31",netmask="255.255.255.0"

**Message:** 2011-10-19 13:26:12.673Z [admin]: Called - set\_safemode

**Message:** 2011-10-19 13:28:12.302Z [admin]: Called - set\_prompt: Args prompt=ip  
**Message:** 2011-10-19 13:30:22.221Z [admin]: Called - set BGP: Args log\_level="debugging"  
**Message:** 2011-10-19 13:31:20.142Z [admin]: Called - set OSPF: Args log\_level="informational"  
**Message:** 2011-10-19 13:32:10.319Z [admin]: Called - set\_nosafemode  
**Message:** 2011-10-19 13:38:42.998Z [admin]: Called - set\_network: Args ip\_address="10.120.20.34 from 10.120.20.31",netmask="255.255.255.0 from 255.255.255.0",gateway\_address="10.120.20.1 from 10.120.20.1"  
**Message:** 2011-10-19 13:41:56.178Z [admin]: Called - set\_ip\_rate\_limit: Args ip\_rate\_limit="on from off"  
**Message:** 2011-10-19 13:43:42.828Z [admin]: Called - set\_monitor\_dns\_alert: Args dns\_alert="on from off"  
**Message:** 2011-10-19 13:46:34.647Z [admin]: updated physical node 0  
**Message:** 2011-10-19 13:46:34.648Z [admin]: Called - set\_interface: Args interface="LAN", speed="100M", duplex="half"  
**Message:** 2011-10-19 13:48:03.066Z [admin]: Called - set\_dns: Args dns="flush all "  
**Message:** 2011-10-19 13:49:35.527Z [admin]: Called - set\_debug: Args all="on from off"  
**Message:** 2011-10-19 09:53:53.595Z [admin]: Called - set\_ibtrap: Args ibtrap="DNS", snmp="true", email="true"  
**Message:** 2011-10-19 09:57:00.747Z [admin]: Called - set\_thresholdtrap: Args thresholdtrap="CpuUsage", trigger="60", reset="50"  
**Message:** 2011-10-19 10:32:50.183Z [admin]: Called - set\_maintenancemode: Args maintenancemode="on from off"  
**Message:** 2011-10-19 14:05:20.132Z [admin]: Called - set\_dhcp\_expert\_mode: Args dhcp\_expert\_mode="true from false"  
**Message:** 2011-10-19 14:07:02.082Z [admin]: Called - set\_dhcp\_release\_delay: Args delay\_time=40 secs  
**Message:** 2011-10-19 14:09:24.285Z [admin]: Called - set\_gsstsig\_key\_expiration\_time: Args gsstsig\_key\_expiration\_time="3000 from 3600"  
**Message:** 2011-10-19 14:10:19.906Z [admin]: Called - set\_named\_worker\_threads: Args named\_worker\_threads="20 from 0"  
**Message:** 2011-10-19 14:11:04.731Z [admin]: Called set\_recursion\_log\_interval: Args recursion\_log\_interval="60"  
**Message:** 2011-10-19 14:14:12.170Z [admin]: Called - set\_partial\_replication: Args partial\_replication="off from on"  
**Message:** 2011-10-19 14:15:33.978Z [admin]: Called - set\_rep\_queue\_ixfr\_limit: Args rep\_queue\_ixfr\_limit="60 from 1000"  
**Message:** 2011-10-19 14:16:16.797Z [admin]: Called - set\_watchdog: Args watchdog\_enabled="true from false"  
**Message:** 2011-10-19 14:17:14.605Z [admin]: Called - set\_fsck  
**Message:** 2011-10-19 14:19:25.282Z [admin]: Called - set\_host\_consistency\_check: Args host\_consistency\_check="on from off"  
**Message:** 2011-10-19 14:21:00.202Z [admin]: Called - set\_internal\_apache\_http\_port: Args internal\_apache\_http\_port="2000 from 9000"  
**Message:** 2011-10-19 14:22:18.682Z [admin]: Called - set\_internal\_jetty\_http\_port: Args internal\_apache\_http\_port="6060 from 8080"  
**Message:** 2011-10-19 14:25:58.704Z [admin]: Called - set\_always\_ret\_nxdomain\_for\_fmz\_ptr: Args always\_ret\_nxdomain\_for\_fmz\_ptr="true from false"  
**Message:** 2011-10-19 14:28:18.046Z [admin]: Called - set\_debug\_tools: Args debug\_tools="db\_binary\_dump"  
**Message:** 2011-10-19 14:29:06.511Z [admin]: Called - set\_dns\_autogen: Args dns\_auto\_gen="check"  
**Message:** 2011-10-19 14:30:54.628Z [admin]: Called - set\_named\_recv\_sock\_buf\_size: Args udp\_so\_rcvbuf="122 from (null)"

## CLI Top Level Commands

**Message:** 2011-10-19 10:33:29.664Z [admin]: Called - delete\_cores\_all  
**Message:** 2011-10-19 10:38:12.356Z [admin]: Called - delete\_cores: Args filename="core.8295.gz"  
**Message:** 2011-10-19 10:58:28.064Z [admin]: Called - delete\_backup\_all  
**Message:** 2011-10-19 11:00:17.917Z [admin]: Called - delete\_backup: Args filename="BACKUP\_6.bkp"  
**Message:** 2011-10-19 12:41:47.707Z [admin]: Called - rotate\_log: Args log="syslog"  
**Message:** 2011-10-19 12:58:11.738Z [admin]: Called - rotate\_log: Args log="audit"  
**Message:** 2011-10-19 12:58:11.738Z [USER\040admin]: rotated the previous audit log to audit.log.0.gz  
**Message:** 2011-10-19 13:51:36.982Z [admin]: Called - reset\_database  
**Message:** 2011-10-19 13:54:14.023Z [admin]: Called - debug\_webui\_restart  
**Message:** 2011-10-19 13:57:39.407Z [USER\040admin]: rebooted the system  
**Message:** 2011-10-19 14:03:41.124Z [admin]: Called - delete\_file: Args groupname="bloxtools", filename="/storage/web-portal/udata/logs/access.log"

## CLI Emergency Commands

**Message:** 2011-10-19 14:32:31.927Z [Emergency\040User]: Called - set\_safemode  
**Message:** 2011-10-19 14:33:23.591Z [Emergency\040User]: Called - set\_nosafemode  
**Message:** 2011-10-19 14:33:41.286Z [Emergency\040User]: Called set\_repsafe\_mode: Args repsafe\_mode = on  
**Message:** 2011-10-19 14:34:47.321Z [Emergency\040User]: Called - set\_weak  
**Message:** 2011-10-19 14:35:25.969Z [Emergency\040User]: Called - set\_fsck  
**Message:** 2011-10-19 14:35:46.604Z [Emergency\040User]: Called - set\_watchdog: Args watchdog\_enabled="true from true"  
**Message:** 2011-10-19 14:41:13.727Z [Emergency\040User]: Called - reset\_database

### Note

During the boot time, if you erroneously press the Enter key before being prompted, NIOS does not wait for a specified time to enter into emergency mode and restarts immediately.

## WAPI Detailed

You can view detailed WAPI session information logs in the audit log for successful WAPI calls such as PUT, POST, and DELETE. For more information, see [Monitoring Tools](#).

**Event:** Member restart or reboot service

**Message:** [2018-07-10 16:23:08.112Z] [admin]: Called(POST) v2.9/member {"\_function": "restartservices", "restart\_option": "FORCE\_RESTART", "service\_option": "ALL"} 3.081 MemberRestartServices: Args service\_option="ALL", grid\_member=Member:infoblox.localdomain, restart\_option="FORCE\_RESTART"

**Event:** All succeeded function calls

**Message:** [2018-07-28 08:56:44.399Z] [admin]: Called(POST) v2.9/network {"\_function": "next\_available\_ip"} 0.034 NextAvailableIp: Args parent=Network:2.2.2.0/24\054network\_view\075default

**Event:** Enhanced audit log for POST method

**Message:** [2018-05-29 09:20:12.026Z] [admin]: Created(POST) v2.9/zone\_auth {"fqdn": "foo.com"} 2.233 AuthZone foo.com DnsView=default: Set fqdn="foo.com"

**Event:** Enhanced audit log for PUT method

**Message:** 2018-06-07 08:45:25.681Z [admin]: Modified(PUT) v2.2/zone\_auth {"comment": "testing auditlogs"} 1.930 AuthZone foo.com DnsView=default: Changed comment: NULL->"testing auditlogs"

**Event:** Enhanced audit log for DELETE method:

**Message:** 2018-07-24 13:11:26.614Z [admin]: Deleted(DELETE) v2.6/zone\_auth {} 0.356 AuthZone foo.com DnsView=default exclude\_subobj=False

## Host Record Logging

NIOS inserts two records for each host record object and the audit log displays the URI, InData and response time twice, that is, one for the host record and the other one for the host address/host alias records.

**Example of Host Record logging:** curl -H "Content-Type: application/json" -k -u admin:infoblox -X POST https://10.120.20.129/wapi/v2.0/record:host -d '{ "ipv4addrs": [ {"ipv4addr": "1.1.1.0", "configure\_for\_dhcp": false, "mac": "aa:0:0:0:1:cc" } ], "comment": "this is my one.perfusera comment", "view": "default", "name": "perfusera.test.com" }'

**Message:** 2018-07-24 12:27:40.375Z [admin]: Created(POST) v2.0/record:host {"ipv4addrs": [ {"ipv4addr": "1.1.1.0", "configure\_for\_dhcp": false, "mac": "aa:0:0:0:1:cc" } ], "comment": "this is my one.perfusera comment", "view": "default", "name": "perfusera.test.com"} 0.236 HostAddress 1.1.1.0 network\_view=default: Set address="1.1.1.0",configure\_for\_dhcp=False,mac\_address="aa:0:0:0:1:cc",match\_option="MAC\_ADDRESS",parent=HostRecord:.\_default.com.foo.perfusera

**Message:** 2018-07-24 12:27:40.375Z [admin]: Created(POST) v2.0/record:host {"ipv4addrs": [ {"ipv4addr": "1.1.1.0", "configure\_for\_dhcp": false, "mac": "aa:0:0:0:1:cc" } ], "comment": "this is my one.perfusera comment", "view": "default", "name": "perfusera.test.com"} 0.236 HostRecord perfusera.foo.com DnsView=default address=1.1.1.0: Set addresses=[address="1.1.1.0"],comment="this is my one.perfusera comment",fqdn="perfusera.foo.com",view=DnsView:default

## Requesting an Object

Each WAPI call for a request object shows the timestamp, user, operation, URI, InData, and the response time.

**Example of Request object:** https://10.35.120.1/wapi/v2.9/request body : { { "method": "POST", "object": "network", "data": { "network": "22.2.2.0/24" } }, { "method": "POST", "object": "network", "data": { "network": "111.1.111.0/24" } } }

**Message:** 2018-10-24 11:18:18.828Z [admin]: Created(POST) v2.9/request [{"object": 'network', 'data': {'network': '22.2.2.0/24'}}, {'method': 'POST'}, {'object': 'network', 'data': {'network': '111.1.111.0/24'}}, {'method': 'POST'}] 5.5867

**Message:** 2018-10-24 11:18:18.828Z [admin]: Created Network 22.2.2.0/24 network\_view=default: Set address="22.2.2.0",cidr=24

**Message:** 2018-10-24 11:18:18.828Z [admin]: Created Network 111.1.111.0/24 network\_view=default: Set address="111.1.111.0",cidr=24

## Scheduling an Object

For a schedule object, PUT/POST/DELETE calls and WAPI session log information, such as URI, InData, and response time, are added only in the first line.

**Example of Schedule object:** curl -k1 -u admin:infoblox -X POST https://10.35.120.1/wapi/v2.9/network -d network=3.3.8.0/24 -d \_schedinfo.scheduled\_time=1540386870

**Message:** 2018-10-24 11:22:01.998Z [admin]: Sched:3 Created(POST) v2.9/network {'\_schedinfo.scheduled\_time': '1540380251', 'network': '3.3.8.0/24'} 1.7615 Network 3.3.8.0/24 network\_view=default: Set address="3.3.8.0",cidr=24

**Message:** 2018-10-24 11:22:01.998Z [admin]: Sched:3 Created ScheduledTask 3: Set scheduled\_time=2018-10-24 11:24:11.000Z,submit\_time=2018-10-24 11:22:01.983Z,submitter="admin",type="SCHEDULED

## Database Backup

NIOS logs information about who started the database backup and where the database backup file is stored.

**Event:** Successful database backup

**Message:** 2020-06-03 10:15:28.634Z [admin]: Called - GetGridData message=backed\040up\040database\040at\040scp\072\root:\*\*\*\*@10.120.20.38/tmp/adf: Args message="backed up database at scp://root:\*\*\*\*@10.120.20.38/tmp/adf" 2020-06-03 10:15:28.852Z [admin]: Called - DataGetComplete message=data\040get\040completed: Args message="data get completed

## Backing Up and Restoring the Database

You must log in with a superuser account to back up files. The administrator must back up system files to the local appliance.

You can restore a backup up file to an appliance running the same NIOS version as that of the appliance from which the

backup file originates. You can also restore a backup file from an appliance running a NIOS version to an appliance running a later NIOS version as long as the upgrade from the earlier NIOS version to the later version is supported. Note that if you need to restore a backup file to an appliance, ensure that the backup file that you are restoring is from an appliance that was Common Criteria compliant as well. NIOS stores information such as who started the database backup and where the backup file is stored in the audit log. For more information see [Audit Log](#).

For more information about backing up and restoring the database, see [Backing Up and Restoring Configuration Files](#).

## About DNS

The TOE provides DNS service. There are two basic methods used to protect DNS communication: TSIG and GSS-TSIG. The TSIG (transaction signature) method signs communications using either HMAC-MD5 or HMAC-SHA25. Both end points must be configured with the key. The GSS-TSIG method (based on the GSS API) uses a Kerberos server to retrieve the key, and is only available in Microsoft environments.

When you configure the TOE to use TSIG and GSS-TSIG keys, you must select HMAC-SHA256 as the key algorithm. For information about using TSIG keys to ensure security in several DNS operations, see the following:

- To control access to DNS views. For more information, see [Defining Match Clients Lists](#) and [Configuration Example: Configuring a DNS View](#).
- To control to which recursive and non-recursive queriers the TOE is allowed to respond. For more information, see [Specifying Queries](#) and [Enabling Recursion](#).
- To authenticate zone transfer requests and replies. For more information, see [Configuring Zone Transfers](#).
- To authenticate and verify dynamic DNS updates from DHCP servers. For more information, see [Enabling DNS Servers to Accept DDNS Updates](#).
- When a secondary DNS server receives DDNS updates, it must forward the updates to the primary server because it cannot update zone data itself. To specify the source of DDNS updates. For more information, see [Forwarding Updates](#).

For information about using GSS-TSIG, see [About GSS-TSIG](#).

## WebUI Settings

This section describes the properties that you can set to ensure the security of the Grid Manager web interface.

### Creating a Login Banner

Before establishing a user session via the WebUI, the TOE displays an initial banner regarding unauthorized use. The message is displayed before the session is established. You can change this message to your organization's specific advisory notice and warning message regarding unauthorized use of the system. For information about defining the login banner, see [Creating a Login Banner](#).

### Modifying the Session Timeout Setting

You can set the length of idle time before an administrative session to the WebUI times out. The default timeout value is 600 seconds (10 minutes). If an admin does not interact with the application for the specified time, the TOE displays a message that a timeout has occurred. The admin is then required to log back in to Grid Manager. For information about setting the session timeout, see [Modifying the Session Timeout Setting](#).

## Managing Certificates

The TOE generates a self-signed certificate when it first starts. Because the default certificate is self-signed, your browser does not have a trusted CA certificate or a cached NIOS appliance server certificate (saved from an earlier connection) to authenticate the NIOS appliance certificate. Also, the hostname in the default certificate



is] www.infoblox.com,[http://www.infoblox.com/] which is unlikely to match the hostname of your NIOS appliance. Consequently, a message appears warning that the certificate is not from a trusted certifying authority and that the hostname on the certificate is either invalid or does not match the name of the site that sent the certificate. To eliminate certificate warnings, you can replace the default self-signed certificate with a different certificate.

After the initial login, you can do one of the following:

- Generate another self-signed certificate with the correct hostname and save it to the certificate store of your browser.
- Request a CA-signed certificate with the correct hostname by generating a Certificate Signing Request (CSR) and sending it to your trusted Certificate Authority (CA). Then when you receive the certificate from the CA, import it to the appliance.

For information about these tasks, see [Managing Certificates](#).

For Common Criteria compliance, superusers must not use CSRs or certificates with keys smaller than 2048 bits. Limited access users are not allowed to upload a certificate with a key that is smaller than 2048 bits, or create a certificate signing request or self-signed certificate with a key size that is smaller than 2048 bits.

## Licenses and Services

The TOE does not have general computing capabilities, other than the services required for the operation, administration and support of the TOE. In the evaluated configuration, the TOE has only the following licenses installed: DNS, DHCP, Microsoft Management, DNS Cache Acceleration and Query Redirection. It does not have the Multi-Grid Management license installed. For more information about licenses, see [Managing Licenses](#).

The following services are disabled by default in the Common Criteria evaluated configuration and no claims are made regarding their function:

- bloxTools
- MGM (Multi-Grid Management)
- HSM (Hardware Security Module) Signing
- Support access
- Remote console access
- Remote authentication using RADIUS and TACACS+

Installing additional licenses or enabling any of the listed services may result in a non-compliant system.

## Enabling / Disabling Common Criteria Mode



Infoblox recommends that you do not change the Common Criteria setting of a NIOS appliance that is in a production environment.

Before you enable Common Criteria mode, you must reset a NIOS appliance to its original factory settings. This removes the database, network settings, logs, and configuration files. Then, it reboots with its factory settings, which are the default user name and password, and default network settings. If you do not reset the appliance to its original factory settings, the appliance will not be Common Criteria compliant, even if you enable Common Criteria mode.

To reset the NIOS appliance to its factory settings:

1. Log in to the Infoblox CLI using a superuser account.
2. Enter the following CLI command:

```
reset all
```

You can enable and disable Common Criteria mode from the Infoblox CLI only. Do the following to set Common Criteria mode on the appliance:

1. Log in to the Infoblox CLI. After executing the **resetall** command, you can log in to the TOE only by using the default superuser admin name **admin** and password **infoblox**.

2. Type the following command:

```
set cc_mode
```

The TOE reboots and goes through boot time self tests. If the test fails, the TOE goes into a loop and displays an error message on the serial console and the LCD. Otherwise, it displays the Login prompt after the self tests.

To clear Common Criteria mode on an appliance, log in to the Infoblox CLI and execute the command: `reset all`.

## Using the CLI

Only superusers can access the CLI. To ensure security, access to the CLI is permitted through a direct console connection only. Note that activating the option **Enable Remote Console Access** in the Grid or Member Properties editor will result in a non-compliant system.

To access the Infoblox CLI through the console port:

1. Connect a serial cable from the console port on your management system to the console port on the appliance. The appliance has a male DB-9 console port on its front panel.
2. Use the following connection settings to launch an emulation session through a serial terminal emulation program such as Hilgraeve Hyperterminal® (provided with the Windows® operating systems):
  - Bits per second: 9600
  - Data bits: 8
  - Parity: None
  - Stop bits: 1
  - Flow control: Xon/Xoff
3. Use the following default user name and password to log in to the Infoblox appliance:

```
admin
```

```
infoblox
```



### Note

After you log in, change the default user name and password of the default superuser admin to prevent unauthorized access to the TOE. For more information on changing passwords, see [Changing the Password and Email Address](#).

For more information about the Infoblox CLI, refer to the *Infoblox CLI Guide*.

## Administration

A user must have an admin account to log in to the TOE. Each admin account belongs to an admin group, which contains roles and permissions that determine the tasks a user can perform.

The TOE provides a default superuser admin group, called **admin-group**, with one superuser administrator, **admin**. The default superuser admin can log in to the TOE, using the default user name **admin** and password **infoblox**.

Superuser admins are the security admins and have full access and control of all the operations of a TOE. Note that you must change the default user name and password of the default superuser admin to prevent unauthorized access to the TOE.

Only superusers can do the following:

- Create admin accounts and groups. For more information, see [Managing Administrators](#).
- Set password parameters. For more information, see [Managing Passwords](#).
- Create the login banner. For more information, see [Creating a Login Banner](#).
- Set the session timeout. For more information, see [Modifying the Session Timeout Setting](#).

Limited-access admin groups provide their members with read-only or read/write access to specific resources. These admin groups can access the appliance through the GUI, API, or both. They cannot access the appliance through the console. In addition, limited-access admins are not allowed to perform the following tasks:

- Download the support bundle.
- Enable SNMP on the appliance.
- Upload files that are larger than 100 MB. If the file size is greater than the maximum size allowed, the **Upload** dialog box closes and an error message is displayed in the feedback panel. The attempt to upload a file that exceeded the maximum will be logged to syslog. Non-superusers only are able to upload files for file distribution and do CSV import.

## Setting Password Restrictions for Local Admins

All admins are required to enter a username and password when they log in to Grid Manager or the CLI. The password is always obscured when an admin logs in. The TOE defaults to locking out the user after three consecutive failed logins. A superuser must define a password policy that is consistent with the security policy of the organization. The password policy specifies the minimum password length and character types, such as lowercase or uppercase characters, that are allowed in the password. In addition, the policy specifies the number of required character changes from the previous password, whether passwords expire and their duration. Additionally, you can require admins to change their passwords when they first log in or after their passwords are reset. For information about defining the password policy, see [Managing Passwords](#).

Local admins must change their passwords according to the defined password policy. A password can be changed as follows:

- By the local admin in the User Profile page. For more information, see [Changing the Password and Email Address](#).
- By the local admin when a password expires or when the admin first logs in. Note that this applies to logging in to the CLI or WebUI.
- By a superuser admin.

## Security Guidelines

Following are security assumptions to ensure that the TOE is administered in a secure manner after it is delivered:

- The environment ensures the physical security of the TOE, commensurate with its value and the value of the data that it contains.
- Administrators are non-hostile, properly trained and trusted to apply all administrator guidance.
- Administrators will take appropriate measures to prevent unauthorized individuals from accessing the TOE.

## Installation and Configuration

To ensure the security of the installation and configuration of the TOE:

- Administrators must install the appliance according to the procedures in the installation guides.
- The TOE contains an option for upgrading the system. This is available only for security administrators. The security administrator will be able to upgrade to a validated release package only. The security administrator can verify the TOE by the version number included in the file name as well as through the administrative interface before and after the upgrade.  
When upgrading, ensure that the .bin2 file is uploaded, and not the .bin file. Refer to the Release Notes of the NIOS version to which the TOE is upgrading for additional upgrade instructions.
- Users' access to the TOE is controlled by security mechanisms and unauthorized users are denied access to the TOE. For more information, see [Administration](#).
- The TOE provides external authentication mechanisms for remote users using SSL with Active Directory. For more information, see [Authenticating Admins Using Active Directory](#).

## Syslog

NIOS appliances generate syslog messages that you can view through the Syslog viewer and download to a directory on your management station. For more information about syslog, see [Using a Syslog Server](#).

Following are the events that are logged and examples of their corresponding syslog messages:

```
\\
*Establishment/Termination* *of* *an* *HTTPS* *Session*
*Event:* Generation of RSA key failed.
*Message:* "Oct 19 09:15:01 EPBYMINW0065T1 httpd\[2115\]: cryptographic key generation
failed"
\\
*Cluster logout*
*Event:* Cluster logout.
*Message:* "2020-06-03T12:58:13+00:00 daemon infoblox.localdomain INFOBLOX-Grid[]: notice
Cluster logout for node <node_name>, for node clean restart.
\\
*Event:* Session is terminated.
*Message:* "Oct 19 09:15:01 EPBYMINW0065T1 httpd\[2115\]:Session terminated (remote
address: 10.6.11.249)"
\\
*Event:* Failed to establish a session.
*Message:* "Oct 19 08:50:21 EPBYMINW0065T1 httpd\[2314\]: Failed to establish a session
(remote address: 10.6.11.249), error 1115 (SSL
routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed)"
\\
*Event:* Session is established.
*Message:* "Oct 19 08:54:42 EPBYMINW0065T1 httpd\[2314\]: Session has been established
(remote address: 10.6.11.249)"
\\
*Establishment/Termination of a* *TLS* *Session*
*Event:* Generation of RSA key failed.
*Message:* "Oct 19 08:38:08 EPBYMINW0065T1 openvpn\[1415\]: cryptographic key generation
failed"
\\
*Event:* Session has been established.
*Message:* "Oct 19 08:38:08 EPBYMINW0065T1 openvpn\[1552\]: Session has been established
(remote address: 10.6.11.249)"
\\
*Event:* HMAC failure:
*Message:* "Oct 19 08:41:01 EPBYMINW0065T1 openvpn\[1567\]: cryptographic key generation
failed: HMAC"
\\
*Event:* Signing failure (constructed message, it is not trivial to obtain it into the
syslog).
*Message:* "Oct 19 08:45:01 EPBYMINW0065T1 openvpn\[1582\]: cryptographic operation
failed: signature"
\\
*Event:* Encryption failure.
*Message:* "Oct 19 08:46:41 EPBYMINW0065T1 openvpn\[1612\]: cryptographic operation
failed: encryption"
\\
*Event:* Decryption failure.
*Message:* "Oct 19 08:46:41 EPBYMINW0065T1 openvpn\[1612\]: cryptographic operation
failed: decryption"
```

```

\\
*Event:* Session was not established.
*Message:* "Oct 19 08:50:21 EPBYMINW0065T1 openvpn\[1701\]: Failed to establish a session
(remote address: 10.6.11.249), error 1115 (SSL
routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed)"
\\
*Event:* Packet was not verified.
*Message:* "Oct 19 08:55:25 EPBYMINW0065T1 openvpn\[1815\]: Packet verification fails
(remote address: 10.6.11.249)"
*Random* *Number* *Generation* *Process*
\[2011/10/19 10:13:46.282\] (26360 /infoblox/one/bin/ib_prngd_control) : ib_prngd daemon
is not running while CC mode is enabled
\[2011/10/19 10:13:46.324\] (26368 /infoblox/one/bin/ib_prngd) main.c:202 main():
ib_prngd daemon starting up...
\[2011/10/19 10:13:46.700\] (26368 /infoblox/one/bin/ib_prngd) main.c:214 main(): Setting
FIPS mode OK \[2011/10/19 10:13:48.400\] (26368 /infoblox/one/bin/ib_prngd) main.c:214
main(): Setting FIPS mode FAILED \[2011/10/19 10:13:46.700\] (26368 /infoblox/one/bin/
ib_prngd) main.c:125 rename_rnd_dev(): Moving
/dev/random to /dev/random_backup OK
\[2011/10/19 10:13:46.700\] (26368 /infoblox/one/bin/ib_prngd) main.c:127
rename_rnd_dev(): Moving
/dev/urandom to /dev/urandom_backup OK
\[2011/10/19 10:13:46.700\] (26368 /infoblox/one/bin/ib_prngd) main.c:234 main():
Creating FIFO
/dev/ib_random OK
\[2011/10/19 10:13:46.700\] (26368 /infoblox/one/bin/ib_prngd) main.c:158
symlink_rnd_dev(): Symlinking
/dev/random to /dev/ib_random OK
\[2011/10/19 10:13:46.700\] (26368 /infoblox/one/bin/ib_prngd) main.c:160
symlink_rnd_dev(): Symlinking
/dev/urandom to /dev/ib_random OK
\[ TIME NOT KNOWN \] (26368) main.c:signal_handler\{\}: ib_prngd received SIGTERM
signal....exiting. \[ TIME NOT KNOWN \] (26368) main.c:signal_handler\{\}: ib_prngd
received SIGINT signal....exiting.
\\
\[ TIME NOT KNOWN \] (26368) main.c:signal_handler\{\}: ib_prngd received SIGQUIT
signal....exiting.
\[ TIME NOT KNOWN \] (26368) main.c:signal_handler\{\}: ib_prngd received an unknown
signal....exiting. \[2011/10/19 10:13:49.205\] (26368 /infoblox/one/bin/ib_prngd)
main.c:135 rename_rnd_dev(): Renaming
/dev/random back OK
\[2011/10/19 10:13:49.205\] (26368 /infoblox/one/bin/ib_prngd) main.c:141
rename_rnd_dev(): Renaming
/dev/urandom back OK
\[2011/10/19 10:13:49.205\] (26368 /infoblox/one/bin/ib_prngd) main.c:255 main():
Removing custom FIFO
/dev/ib_random OK
\[2011/10/19 10:13:49.205\] (26368 /infoblox/one/bin/ib_prngd) main.c:255 main():
Removing custom FIFO
/dev/ib_random FAILED
\[2011/10/19 10:13:49.205\] (26368 /infoblox/one/bin/ib_prngd) main.c:141
rename_rnd_dev(): Renaming
/dev/urandom back FAILED
\[2011/10/19 10:13:49.205\] (26368 /infoblox/one/bin/ib_prngd) main.c:135
rename_rnd_dev(): Renaming
/dev/random back FAILED
\[2011/10/19 10:25:22.931\] (26557 /infoblox/one/bin/ib_prngd) main.c:189 main(): Error!
/infoblox/one/bin/ib_prngd is already running

```

```

\[2011/10/19 10:26:58.107\] (26560 /infoblox/one/bin/ib_prngd) main.c:52 self_test():
OpenSSL FIPS mode functionality self test OK
\[2011/10/19 10:26:58.107\] (26560 /infoblox/one/bin/ib_prngd) main.c:52 self_test():
OpenSSL FIPS mode functionality self test FAILED
*Failures on Invoking* *Functionality*
*Event:* Invalid size specified for algorithm HMAC-SHA256.
*Message:*{*}2011-10-19T17:57:12-04:00 user EPBYMINW2856 httpd\[ \]: err TSIG key
generation failure: Size 512 can not be used with algorithm HMAC-SHA256
\\
*Event:* Invalid algorithm specified in Common Criteria mode.
*Message:* 2011-10-19T18:12:22-04:00 user EPBYMINW2856 httpd\[ \]: err TSIG key (keylen =
256, alname = HMAC-MD5) generation error : Only HMAC-SHA256 available in CC mode.
*Open VPN*
*Event:* Generation of RSA key failed
*Message:* Oct 19 08:38:08 EPBYMINW0065T1? openvpn\[1415\]: cryptographic key generation
failed
\\
*Event:* Session has been established
*Message:* Oct 19 08:38:08 EPBYMINW0065T1? openvpn\[1552\]: Session has been established
(remote address: 10.6.11.249)
\\
*Event:* HMAC failure
*Message:* Oct 19 08:41:01 EPBYMINW0065T1? openvpn\[1567\]: cryptographic key generation
failed: HMAC
\\
*Event:* Signing failure
*Message:* Oct 19 08:45:01 EPBYMINW0065T1? openvpn\[1582\]: cryptographic operation
failed: signature
\\
*Event:* Encryption failure
*Message:* Oct 19 08:46:41 EPBYMINW0065T1? openvpn\[1612\]: cryptographic operation
failed: encryption *Event:* Decryption failure
\\
*Message:* Oct 19 08:46:41 EPBYMINW0065T1? openvpn\[1612\]: cryptographic operation
failed: decryption
\\
*Event:* Session was not established
*Message:* Oct 19 08:50:21 EPBYMINW0065T1? openvpn\[1701\]: Failed to establish a session
(remote address: 10.6.11.249), error 1115 (SSL
routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed)
\\
*Event:* Packet was not verified
*Message:* Oct 19 08:55:25 EPBYMINW0065T1? openvpn\[1815\]: Packet verification fails
(remote address: 10.6.11.249)
*HTTPS*
*Event:* Generation of RSA key failed
*Message:* Oct 19 09:15:01 EPBYMINW0065T1? httpd\[2115\]: cryptographic key generation
failed
\\
*Event:* Session is terminated
*Message:* Oct 19 09:15:01 EPBYMINW0065T1? httpd\[2115\]: Session terminated (remote
address: 10.6.11.249)
\\
*Event:* Failed to establish a session
*Message:* Oct 19 08:50:21 EPBYMINW0065T1? httpd\[2314\]: Failed to establish a session
(remote address: 10.6.11.249), error 1115 (SSL
routines:SSL3_GET_SERVER_CERTIFICATE:certificate verify failed)
\\

```

```

*Event:* Session is established
*Message:* Oct 19 08:54:42 EPBYMINW0065T1? httpd\[2314\]: Session has been established
(remote address: 10.6.11.249)
\\
*Event:* HMAC failure
*Message:* Oct 19 08:55:56 EPBYMINW0065T1? httpd\[2356\]: cryptographic key generation
failed: HMAC
\\
*DNS*
*Message:* 2011-10-18T13:37:33+00:00 daemon (none) named\[4456\]: err client
10.32.2.108#47160: request has invalid signature: TSIG sha256cc: tsig verify failure
(BADKEY) 2011-10-18T13:37:33+00:00 daemon (none) named\[4456\]: err client
10.32.2.108#47160: request has invalid signature: TSIG sha256cc: tsig verify failure
(BADKEY)
*DHCP*
*Message:* 2011-10-18T11:18:38+00:00 daemon (none) dhcpd\[20440\]: err No tsec for use
with key sha128cc *Message:* 2011-10-31T18:32:17+00:00 daemon (none) dhcpd\[20440\]: err
Invalid operation in ddns code.
*Upgrade*
*Message:* 2011-10-26T12:33:30-04:00 user EPBYMINW2994t1 infoblox_crypt\[\\]: err
cryptographic operation failed: decryption
*Message:* 2011-10-26T12:34:33-04:00 user EPBYMINW2994t1 infoblox_crypt\[\\]: err
cryptographic operation failed: encryption
*Message:* 2011-10-26T12:35:53-04:00 user EPBYMINW2994t1 infoblox_crypt\[\\]: err
cryptographic operation failed: RSA verify signature
*Message:* 2011-10-26T12:38:56-04:00 user EPBYMINW2994t1 infoblox_crypt\[\\]: err
cryptographic operation failed: RSA signing
\\
*Quotas*
*Event:* When the administration backend is overloaded by too much combined GUI and API
traffic, a message like this is logged to syslog (it is not associated with any user).
*Message:* 2011-10-31T23:42:21+00:00 user (none) httpd\[\\]: warning Too many
administration connections *Event:* Disk space limit was changed and is below the disk
usage.
*Message:* 2011-11-02T00:24:54+00:00 user manojk-vm httpd\[\\]: err Storage Limit has been
lowered and usage now exceeds the limit, Usage: 150 MB, Limit :100 MB
*Event:* Disk space limit reached.
*Message:* 2011-11-02T00:24:54+00:00 user manojk-vm httpd\[\\]: err Exceed the TFTP
Storage limit, User name:user1, Used Storage:2048 B, File name :a.zip, File size
:272629904 B, Limit :102400 B
*Open* *SSL*
*Event:* FIPS self test failed.
*Message:* FIPS routines:EVP_DigestInit_ex:fips selftest failed:digest.c:18: *Event:*
Tried to use non-FIPS algorithm in FIPS mode.
*Message:* 140576691959464:error:140A9129:SSL routines:SSL_CTX_new:only tls allowed in
fips mode:ssl_lib.c:1527:
*Message:* 139852903503528:error:0A07C06E:dsa
routines:func(124):reason(110):dsa_key.c:131: *Event:* Used DES-CBC-SHA cipher suite in
FIPS mode.
*Message:* 140418599392936:error:1410D0B9:SSL routines:SSL_CTX_set_cipher_list:no cipher
match:ssl_lib.c:1282:
*Event:* Error setting digest MD5.
*Message:* 140403566474920:error:060800A0:digital envelope
routines:EVP_DigestInit_ex:unknown cipher:digest.c:248:
*Replay* *Detection*
*Event:* OpenVPN
*Message:* Mon Oct 22 22:30:00 2007 us=939054 Authenticate/Decrypt packet error: bad
packet ID (may be a replay): \[ #0 / time = (4196958004) Wed Nov 23 16:11:48 1966 \]

```

```

silence this warning with --mute-replay-warnings, error_prefix, packet_id_net_print
(&pin, true, &gc)
*Event:* OpenVPN
*Message:* Mon Oct 22 22:30:00 2007 ACK reliable_can_send is a replay : \[1\] 0 *Event:*
HTTPS
*Message:* Mon Oct 22 22:30:00 2007 Digest: Warning possible replay attack: nonce-count
check failed: 12345678
= 123456789
\\
*GSS-TSIG*
*Message:* 2011-10-18T13:37:33+00:00 named\[4456\]: err signature invalid: message
integrity
*Message:* 2011-10-18T14:32:22+00:00 named\[4456\]: err authentication failed for aes128-
cts-hmac-sha1-96: unknown principal
*Message:* 2011-10-18T14:42:12+00:00 named\[4456\]: err signature failed to verify(1)
*Message:* 2011-10-18T14:45:54+00:00 named\[4456\]: err signature is in the future
*User* *Login*
*Message:* 2011-10-19T08:27:23-04:00 user spradhan-vm serial_console\[ \]: info User admin
set_repsafe_mode: On *Message:* 2011-10-19T08:29:54-04:00 user spradhan-vm
serial_console\[ \]: info User admin set_repsafe_mode: Off
*Message:* 2011-10-19T08:38:02-04:00 user spradhan-vm serial_console\[ \]: info audit has
been truncated to approximately 2011-10-19T08:29:00-04:00
\\
*Message:* 2011-10-19T08:41:47-04:00 user spradhan-vm serial_console\[ \]: info syslog has
been truncated to approximately 2011-10-19T08:41:00-04:00
*File* *Rotation*
*Event:* Audit log is rotated.
*Message:* 2011-11-01T18:23:00-07:00 user manojk-vm perl\[18990\]:info audit has been
truncated to approximately 2011-11-01T18:23:00-07:00
*Event:* Syslog is rotated.
*Message:* 2011-11-01T18:23:00-07:00 user manojk-vm perl\[18990\]:info syslog has been
truncated to approximately 2011-11-01T18:23:00-07:00
*Zeroization*
*Event:* Logged in case of error
*Message:* 2011-11-01T15:32:59-04:00 daemon manojk-vm ntpd\[18990\]:err Error erasing /
storage/etc/ntp.keys using shred
*First* *Login*
*Message:* \[2011/10/19 08:44:45.866\] (32289 /usr/bin/httpd)
/infoblox/common/lib/python/infoblox/one/admin_conn/userauth.py:415 _log(): \[user\]
First_Login to=AdminConnector auth=LOCAL group=admin-group apparently_via=GUI
*Password* *Expired*
*Message:* \[2011/10/20 09:17:29.257\] (15750 /usr/bin/httpd)
/infoblox/common/lib/python/infoblox/one/admin_conn/userauth.py:415 _log(): \[user\]
Password_Expired to=AdminConnector ip=127.0.0.1 auth=LOCAL group=admin-group
apparently_via=GUI
*Password* *Reset*
*Message:* \[2011/10/19 08:44:45.962\] (32289 /usr/bin/httpd)
/infoblox/common/lib/python/infoblox/one/admin_conn/userauth.py:415 _log(): \[user\]
Password_Reset to=AdminConnector auth=LOCAL group=admin-group apparently_via=GUI
*Failed* *Password* *Reset*
*Message:* \[2011/10/19 09:07:33.343\] (32526 /usr/bin/httpd)
/infoblox/common/lib/python/infoblox/one/admin_conn/userauth.py:415 _log(): \[user\]
Password_Reset_Error to=AdminConnector auth=LOCAL group=admin-group apparently_via=GUI
\\
\\
\\
\\

```



# CSV Import Reference

This section provides reference information during a CSV import. It contains the following topics:

- [Documentation Conventions](#)
- [Customer Care](#)
- [Guidelines for CSV Import](#)
- [CSV File Format](#)
- [Supported Object Types](#)
- [Importing Multiple Action CSV file](#)

## Documentation Conventions

The text in this guide uses the following style conventions:

Style	Usage
<b>bold</b>	Indicates examples of the field names.
<code>data</code>	Signifies the data in a CSV file.

## Customer Care

This section addresses user accounts, software upgrades, licenses and warranties, and technical support.

### User Accounts

The Infoblox appliance ships with a default user name and password. Change the default admin account password immediately after the system is installed to safeguard its use. Make sure that the NIOS appliance has at least one administrator account with superuser privileges at all times, and keep a record of your account information in a safe place. If you lose the `admin` account password, and did not already create another superuser account, the system will need to be reset to factory defaults, causing you to lose all existing data on the NIOS appliance. You can create new administrator accounts, with or without superuser privileges. For more information, refer to the *Infoblox Administrator Guide*.

### Software Upgrades

Software upgrades are available according to the Terms of Sale for your system. Infoblox notifies you when an upgrade is available. Register immediately with Infoblox Technical Support at <http://www.infoblox.com/support/customer/evaluation-and-registration> to maximize your Technical Support.

### Technical Support

Infoblox Technical Support provides assistance via the Web, e-mail, and telephone. The Infoblox Support web site at <https://support.infoblox.com> provides access to product documentation and release notes, but requires the user ID and password you receive when you register your product online at: <http://www.infoblox.com/support/customer/evaluation-and-registration>.

## Guidelines for CSV Import

This section provides general guidelines and file format information about each supported object type for CSV import. You must follow the format and syntax described in this section to ensure a successful data import.

You can create a data file using a text editor, such as Microsoft Notepad, or an application that supports CSV file format, such as Microsoft Excel. You can also import data using Infoblox Migration Wizard, which is a standalone software tool that facilitates the migration of DNS and DHCP data from Microsoft servers to the Infoblox Grid. This tool synchronizes DNS and DHCP data from Microsoft servers and generates a CSV file based on conversion rules you set up through the tool. You can then import the CSV data to the Infoblox Grid through CSV import. For more information, refer to the *Infoblox Administrator Guide for Infoblox Migration Wizard*.



CSV imports and operations that involve massive data, such as deleting large zones and recursive deletion of networks and all child objects, will significantly affect member performance, resulting in service outage.

### General Guidelines

Follow these rules to start a data file:

- Do not use UTF-8 characters in the CSV file name, but the contents of a CSV file must be encoded in UTF-8 characters. Note that Microsoft Excel imports data in the default code page, either in ISO-8859-1 or WINDOWS-1252. You must not import a CSV file that is encoded in Windows 1252 or ISO-8859-1 formats.
- Use a new line to enter data for each row. Separate each data field with a supported separator, such as a comma, semicolon, space, or tab.
- Do not include blank lines in the data file.
- Field names: Specify the field names in the second line. You can include multiple rows of field names as long as you define the fields before the data. The first column in the field name row must be defined as "Header." The rest of the columns are field names of the data. Columns without a field name are ignored. If multiple field names are specified, the latest field names are used to import the data.
- It may take longer than expected to import a large number of DHCP ranges that are associated with a single MAC address filter.
- When a CSV import starts, the appliance validates the first 100,000 rows of data in the CSV file. If the file contains more than 100,000 rows of data, the appliance validates the rest of the data as the import progresses.
- The appliance supports up to one million rows of data in each CSV import.
- Use the add function to add new rows from the imported CSV file to the database.
- Use the **Override** function, not the merge function, to overwrite existing data. When you use the merge function, the appliance does not overwrite existing data, even if the data file contains new data.
- Use the **Delete** function to delete import jobs that are uploaded. You can delete the content of a CSV file that you have imported to the database. Note that you cannot delete jobs that are already imported.
- Use the **Replace** operation to replace current data in the database with data in the imported file. You can use the replace function for authoritative zone data only. The replace operation might affect system performance if you try to replace a zone with a lot of changes. Infoblox recommends that you perform the replace operation for large import files (more than 10,000 rows of changes) during non-peak hours. This operation does not support DNS records that are automatically generated or exported, but it supports NS records that are created manually.
- When you import CSV files for NS record updates, you must specify a value for **zone\_nameservers**. NIOS displays an error message if you do not specify a value for this field when you import the CSV file.
- When you perform a CSV export of automatically created NS records using Infoblox API, the **zone\_nameservers** field will have an empty value. Therefore, if you import the previously exported CSV file that includes automatically created NS records through the Infoblox GUI, then the CSV import fails and Grid Manager displays an error message.
- When you perform a CSV import that includes objects that have scheduled changes or updates associated with them, the import fails. Only superusers can cancel the scheduled changes.
- When you stop an import, the appliance completes the import of the data row it is currently processing before it stops the import. You cannot resume the import from where it stopped.
- You can download uploaded or error files, snapshots, and results file. For more information, see *Infoblox NIOS Administrator Guide*.

- When you import a small file, the appliance processes the import quickly. Under this circumstance, the appliance may generate an error message when you try to stop the import because the import is completed before you can stop it.
- The error files of the last two imports are stored on the appliance. You can download these files using the `APIimport_id` method. For information, refer to the *Infoblox API Documentation*.
- When you use Microsoft Excel to create or view a data file, ensure that you review the settings of the file. Some data, such as dates, may show up in a different format depending on your settings.
- All operations triggered by a CSV import are recorded in the audit log.

## Data Specific Guidelines

Follow these guidelines to enter data:

- The appliance uses double quotes (") as the escape characters in CSV import. If you want to include supported separators in a field, you must enclose the data in a pair of double quotes (" "). This applies to the field names and data. For example, if you want to use the field name **ADMGRP-CSV,;Import**, you must enter **"ADMGRP-CSV,;Import"** as the value. Otherwise, the import fails. When you enter "123""123", the imported data is 123""123, and when you enter "\\", the imported data is \.
- If you have an empty value in the last field, you must still include the separator at the end of the data row. Otherwise, the corresponding column and all its data are not imported, and the appliance generates an error.
- For each supported object type, you must include all the required fields in the data file. For information, see [Supported Object Types](#). All required fields are marked with an asterisk (\*) in an exported file. Note that if you want to modify a required field XXX (for either the overwrite or merge function), you must add a corresponding field, *new\_XXX*, to include the new value. For example, "fqdn" is a required field in an A record. If you want to update this field, you must include a new field "\_new\_fqdn" and define the new value here. The appliance overwrites the existing data in the required field using the values you specify in the new field. Note that the replace function ignores \_\_new\_XXX fields in the imported CSV files.
- When you perform an overwrite function, you must define all boolean and integer data types in each supported object type in order for the appliance to overwrite existing data.
- Inherited fields: The appliance uses the following conventions to override inherited fields:
  - When a value is specified in a field, the appliance overrides the inherited value with the new value.
  - When a value is set to " " or an empty value, the appliance does not override the inherited value.
  - When a value is set to a string with a value of <empty>, the appliance overrides the inherited value with an empty value.
- Extensible attributes: A field name of EA-XXX indicates an extensible attribute, where XXX is the attribute name. The value of an extensible attribute can be a string, a list, an integer, an email address, a URL, or a date in YYYY-MM-DD format. Note that extensible attributes do not support time zones. Following are some examples:
  - "EA-Site" is a predefined string type for locations. It can have a value of "Santa Clara".
  - "EA-User" is a user defined list type for employee types. It can have a list of values, such as "Local,Remote,Temp". Note that only one value can be specified when importing the extensible attribute.
  - "EA-Building" is a predefined integer type for building numbers. It can have a value of "5".
  - "EA-TechPubs" is a user defined email address type. It can have a value of "techpubs@infoblox.com".
  - "EA-IB" is a user defined URL type. It can have a value of "[www.infoblox.com](http://www.infoblox.com)".
  - "EA-Date" is a user defined date type attribute. It can have a value of "2010-11-20".
- Admin permissions: A field name of ADMGRP-XXX indicates the admin permission of a specific admin group, where XXX is the name of the admin group. The value of an admin permission can be a string or a list of strings with subtypes. If there is a single value in the permission, use RW, RO, or DENY. If there is a subtype in the permission, use a list format, such as "RW, ARecord/RO".
- DHCP options: A field name of OPTION-XXX-nn indicates a DHCP option, where XXX is the vendor name of the option and nn is the option number. If the option is of the DHCP vendor class, you can omit -XXX in the field name. For example, OPTION-1 implies vendor class = DHCP and option number = 1, and OPTION-CISCO-122 implies vendor class = CISCO and option number = 122.
- Named ACLs (access control lists): When you import a named ACL or ACEs (access control entries) to a named ACL, ensure that you validate the named ACLs to avoid conflicts and unexpected results. When adding ACEs to a named ACL, all entries are appended to the end of the list. To reorder ACEs in a named ACL through CSV import, you must first export the ACEs, delete all the ACEs in the current ACL, reorder the ACEs in the exported .csv file, and then re-import the ACEs to the named ACL. For more information about access control and named ACLs, refer to the *Infoblox NIOS Administrator Guide*.

## CSV Import Limitations

Ensure that you understand the following limitations before you start an import:

- You can import multiple CSV files at a time, but at any given time you can execute only one single task. The import tasks are queued. Note that only one task at a time will be in the **Import in progress** state, while the others are in the **Import pending** state.
- You cannot roll back to previous data.
- The following data cannot be imported: Microsoft management, DNSSEC, and GSS-TSIG data.
- You cannot export or import zone configuration for DNSSEC signed zones, although resource records added for a signed zone are supported.
- Only editable data can be imported. Discovered data cannot be imported or manipulated.
- If you upload a file and preview the file using the **Preview** option, and later update the content of the same CSV file, and then try to view the edited file using the same *Preview* wizard, you may not be able to see the changes. Infoblox recommends that you start a fresh CSV import to upload the edited file and navigate to the *Preview* wizard to preview the file.
- You cannot perform the CSV import operation on a Microsoft Server zone object, but NIOS allows you to perform the CSV import operation on records within a Microsoft Server zone. You may not see an error message when you perform a CSV import using the replace operation on an Microsoft Server zone.
- When you promote a new Grid Master during an import, the import stops; and it does not restart on the new Grid Master. When a failover occurs during an import, the import stops on the old active node, and it does not restart on the new active node.
- When you configure Unbound as the DNS resolver, NIOS does not support certain features and they are not displayed in Grid Manager. However, fields related to these unsupported features will appear in CSV export files. Although these fields are only relevant to the IB-4030-10GE appliance and might not apply to the appliances in your Grid, you can still perform CSV imports using these CSV export files without any issues.

## Limitations for CSV Import and Export

Ensure that you understand the following limitations before you start an import:

- You can import multiple CSV files at a time, but at any given time you can execute only one single task. The import tasks are queued. Note that only one task at a time will be in the **Import in progress** state, while the others are in the **Import pending** state.
- You cannot roll back to previous data.
- You cannot import network containers.
- The following data cannot be imported: Microsoft management, DNSSEC, and GSS-TSIG data.
- You cannot export or import zone configuration for DNSSEC signed zones, although resource records added for a signed zone are supported.
- Only editable data can be imported. Discovered data cannot be imported or manipulated.
- If you upload a file and preview the file using the **Preview** option, and later update the content of the same CSV file, and then try to view the edited file using the same *Preview* wizard, you may not be able to see the changes. Infoblox recommends that you start a fresh CSV import to upload the edited file and navigate to the *Preview* wizard to preview the file.
- You cannot perform the CSV import operation on a Microsoft Server zone object, but NIOS allows you to perform the CSV import operation on records within a Microsoft Server zone. You may not see an error message when you perform a CSV import using the replace operation on an Microsoft Server zone.
- When you promote a new Grid Master during an import, the import stops; and it does not restart on the new Grid Master. When a failover occurs during an import, the import stops on the old active node, and it does not restart on the new active node.
- When you configure Unbound as the DNS resolver, NIOS does not support certain features and they are not displayed in Grid Manager. However, fields related to these unsupported features will appear in CSV export files. Although these fields are only relevant to the IB-4030-10GE appliance and might not apply to the appliances in your Grid, you can still perform CSV imports using these CSV export files without any issues.
- During the start of the CSV export, only the approximate number of objects to be exported is displayed on Grid Manager. You can get the total exported object count only after the CSV export is complete.

## CSV File Format

A CSV file is typically created and edited using a spreadsheet, though you can create a CSV file in a text editor using any supported separator. You can include more than one object type in a single CSV file when you add or modify data. For information, see [Supported Object Types](#). You can also organize field names and data in a CSV file using different formats, as shown in the following examples. For additional information about how to create a CSV file, see [Guidelines for CSV Import](#).

You can create one CSV file to update data of multiple object types (Network and Host Record), as illustrated in [CSV File Example 1](#). In this example, you define the field names you want to modify for the two object types in rows 1 and 2. You then include the corresponding data as shown in rows 3 to 6.

### CSV File Example 1

	A	B	C	D	E
1	HEADER-NETWORK	ADDRESS*	NETMASK*	EA-Gateway	EA-Secondary Address
2	HEADER-HostRecord	configure_for_dns*	FQDN*	ADDRESSES	
3	NETWORK	10.251.133.128	255.255.255.192	10.251.133.129	
4	NETWORK	10.176.80.255	255.255.252.0	10.176.80.1	172.16.213.0
5	HostRecord	TRUE	host1.dhcp.corp100.com	172.20.2.21	
6	HostRecord	TRUE	host2.dhcp.corp100.com	172.20.2.22	

In the above example the field name HEADER-NETWORK identifies the first row as a header row for the Network objects. The field names ADDRESS, NETMASK, EA-Gateway, and EA-Secondary Address (in rows B1 to E1) tell NIOS how to interpret a row of network data in the CSV file. Each row of data that begins with "Network" in column A is identified as a network data row. Therefore, NIOS interprets rows 3 and 4 as network data rows, in which column B contains the network addresses, column C contains the network masks, and columns D and E contain extensible attribute values for gateway and secondary address.

Similarly, the field name HEADER-HostRecord identifies the second row as a header row for the Host Record objects. This header declaration tells NIOS that for each subsequent row of data that begins with "HostRecord" in column A, column C contains the FQDN of the host, and column D contains the host address. Therefore, NIOS interprets rows 5 and 6 as host record data rows that contain the FQDNs of the hosts in column C and the host addresses in column D. Alternatively, you can organize the information in the table above so that the data rows immediately follow the header rows, as shown in [CSV File Example 2](#).

### CSV File Example 2

	A	B	C	D	E
1	HEADER-NETWORK	ADDRESS*	NETMASK*	EA-Gateway	EA-Secondary Address

2	NETWORK	10.251.133.128	255.255.255.192	10.251.133.129	
3	NETWORK	10.176.80.255	255.255.252.0	10.176.80.1	172.16.213.0
4	<b>HEADER-HostRecord</b>	<b>configure_for_dns*</b>	<b>FQDN*</b>	<b>ADDRESSES</b>	
5	HostRecord	TRUE	host1.dhcp.corp100.com	172.20.2.21	
6	HostRecord	TRUE	host2.dhcp.corp100.com	172.20.2.22	

You can also specify multiple header declarations for the same object type, as shown in [CSV File Example 3](#). In this example, you specify the field names (in row 1) and data (in rows 2 and 3) to modify the extensible attributes (EA-Gateway and EA-Secondary Address) of two network addresses. You then specify the field names (in row 4) and data (in rows 5 and 6) to add new extensible attributes (EA-Gateway and EA-Host Range) of two other network addresses.

### CSV File Example 3

	A	B	C	D	E
1	<b>HEADER-NETWORK</b>	<b>ADDRESS*</b>	<b>NETMASK*</b>	<b>EA-Gateway</b>	<b>EA-Secondary Address</b>
2	NETWORK	10.251.133.128	255.255.255.192	10.251.133.129	
3	NETWORK	10.176.80.255	255.255.252.0	10.176.80.1	172.16.213.0
4	<b>HEADER-NETWORK</b>	<b>ADDRESS*</b>	<b>NETMASK*</b>	<b>EA-Gateway</b>	<b>EA-Host Range</b>
5	NETWORK	10.176.90.0	255.255.255.128	10.176.90.1	10.176.90.4-126
6	NETWORK	10.176.90.128	255.255.255.128	10.176.90.129	10.176.90.132-254

The examples in this section are illustrated using tables that resemble spreadsheet layouts. However, all other examples in this appendix use the comma separated value text file format. For example, the following is the CSV file notation equivalent of [CSV File Example 2](#).

```

HEADER-NETWORK,ADDRESS*,NETMASK*,EA-Gateway,EA-Secondary Address
NETWORK,10.251.133.128,255.255.255.192,10.251.133.129
NETWORK,10.176.80.255,255.255.252.0,10.176.80.1,172.16.213.0
HEADER-HostRecord,FQDN*,ADDRESSES
HostRecord,host1.dhcp.corp100.com,172.20.2.21
HostRecord,host2.dhcp.corp100.com,172.20.2.22

```

## IDN Support for CSV Import

The appliance supports IDNs (Internationalized Domain Names) and punycode for most of the DNS object types in a CSV file. An IDN is a domain name that contains a language-specific script or alphabet, such as Arabic, Chinese, Russian, Devanagari, or the Latin alphabet-based characters with diacritics, such as French. IDNs are encoded in multi-byte Unicode and are decoded into ASCII strings using a standardized mechanism known as Punycode transcription. For example, DNS Zone 'инфоблок.рф' (IDN in Russian) can be written as 'xn-90anhdigczv.xn-p1ai' in the punycode representation. For information about IDNs, refer to the Infoblox NIOS Administrator Guide.

You can use either IDNs or punycode to create DNS zones. Even if you use punycode to create a zone, the appliance automatically generates the corresponding IDN and displays the zone name in its native characters. Note that the appliance does not perform any conversion (IDN to punycode and vice versa) for DNS records, but preserves the data in the original characters. In other words, if a DNS object or a field name contains IDN, the appliance imports the data in IDN. If a DNS object or a field name is in punycode, the appliance imports the data in punycode. For more information about supported objects for CSV import, see [Supported Object Types](#).

## CSV Import for Response Policy Zones

You can import local RPZs (Response Policy Zones) and their rulesets using the **CSV Import** feature. When you import local RPZs using this feature, you must specify three new columns, **priority**, **rpz\_policy**, and **substitute\_name** with relevant values, whereas importing an RPZ ruleset requires specifying the value for parent RPZ in the **parent\_zone** column, as mentioned in the following tables. For a local RPZ, CSV import supports all the values that are listed in [Authoritative Zone](#) along with the three new columns. However, for RPZ rulesets it supports the values that are listed in [CNAME Record](#) along with a new column **parent\_zone**.

For example, if you want to add a new local RPZ, **JKL.INFO** and substitute this domain with **JKI.NET**, then you must mention the priority, rpz\_policy, and substitute name as follows:

A	B	C	D	E	F	G	H	I
HEADER-RESPONSEPOLICYZONE	FQDN*	ZONE_FORMAT*	ALLOW_UPDATE	PRIORITY	RPZ_POLICY	SUBSTITUTE_NAME	VIEW	ZONE_TYPE
RESPONSEPOLICYZONE	ABC.NET	FORWARD	TSIG-RPZ_LOCAL_UPDATER_KEY_default.abc.net/kA36uJeavmhrH2Yqx8hEDPC6okSFcsOb2evyWVAO5fM=/ALLOW/HMAC-SHA256	1001	GIVEN		DEFAULT	RESPONSE POLICY
RESPONSEPOLICYZONE	XYZ.IN	FORWARD	TSIG-RPZ_LOCAL_UPDATER_KEY_default.xyz.in/kA36uJeavmhrH2Yqx8hEDPC6okSFcsOb2evyWVAO5fM=/ALLOW/HMAC-SHA256	1002	NXDOMAIN		DEFAULT	RESPONSE POLICY

A	B	C	D	E	F	G	H	I
RESPONSEPOLICYZONE	AIM.EDU	FORWARD	TSIG-RPZ_LOCAL_UP DATER_KEY._default.aim.edu/ vleLOfean4 YZUMOzGivWnxht OP XWM5QfJwxftHjBD XjQ =/ALLOW/HMAC-SHA 256	1003	NODATA		DEFAULT	RESPONSE POLICY
RESPONSEPOLICYZONE	PQDR.COM	FORWARD	TSIG-RPZ_LOCAL_UP DATER_KEY._default.pqdr.com/ R9TDpx8N +cBs0W32hEDzk5 M wRjPuH%2FeYJsS gUk sX8SM=/ ALLOW/HMAC-SHA256	1004	PASSTRU		DEFAULT	RESPONSE POLICY
RESPONSEPOLICYZONE	JKL.INFO	FORWARD	TSIG-RPZ_LOCAL_UP DATER_KEY._default.jkl.info/ rLopR5+Sf4M pcfYpDJV+KWAdtT XA U5kFTFWFWuLV2 Rw= /ALLOW/HMAC-SHA2 56	1005	SUBSTITUTE	JKI.NET	DEFAULT	RESPONSE POLICY
RESPONSEPOLICYZONE	ASAC.COM	FORWARD	TSIG-RPZ_LOCAL_UP DATER_KEY._default.asac.com/ kA36uJeav mhrH2Yqx8hEDPC 6o kSFcsOb2evyWVA O5 fM=/ALLOW/ HMAC-S HA256	1006	DISABLED		DEFAULT	RESPONSE POLICY

Examples of Substitute and Block Domain Names:

The following example shows a new column, **parent\_zone**, which is added to the spreadsheet while importing an RPZ ruleset to a local RPZ **abc.net**:

A	B	C	D	E	F
HEADER- RESPONSEPOLICYCNAMERECORD	FQDN*	CANONICAL_NAME	DISABLED	PARENT_ZONE	VIEW
RESPONSEPOLICYCNAMERECORD	CLARITY.ABC.NET	CLEAR.IN	FALSE	NET.ABC	DEFAULT
RESPONSEPOLICYCNAMERECORD	ARM.ABC.NET		FALSE	NET.ABC	DEFAULT



Example of an A Record CSV format:

A	B	C	D	E	F
HEADER- RESPONSEPOLICYARECORD	ADDRESS*	FQDN*	DISABLED	PARENT_ZONE	VIEW
RESPONSEPOLICYCNAMERECORD	10.32.2.1	PQR.ABC.NET	FALSE	NET.ABC	DEFAULT

Example of an RPZ Policy IP Address:

A	B	C	D	E	F
HEADER- RESPONSEPOLICYIPADDRESS	FQDN*	CANONICAL_NAME	DISABLED	PARENT_ZONE	VIEW
RESPONSEPOLICYIPADDRESS	10.1.2.3.ABC.NET	10.1.2.3	FALSE	NET.ABC	DEFAULT

Example of an RPZ Policy Client IP Address:

A	B	C	D	E	F
HEADER- RESPONSEPOLICY CLIENTIPADDRESS	FQDN*	CANONICAL_NAME	DISABLED	PARENT_ZONE	VIEW
RESPONSEPOLICY CLIENTIPADDRESS	10.1.2.1.ABC.NET	10.1.2.1	FALSE	NET.ABC	DEFAULT

Note the following:

- You must specify the name of the parent zone when you import RPZ rules to a local zone. For example, **clarity.abc.net** where **abc.net** is the local RPZ.
- In the above example, the domain name **clarity.abc.net** is substituted with the domain name **clear.in** because **clear.in** is specified as the canonical name.
- The domain **arm.abc.net** is blocked and the DNS client receives a message that the domain does not exist. For more information about RPZ rules, refer to the *Infoblox NIOS Administrator Guide*.

## CSV Format for Inheritable Extensible Attributes

### Exporting Inheritable Extensible Attributes

When you export data, if an object has inheritable extensible attributes associated with it, then an additional column **EAIherited-XXX** is displayed in the spreadsheet, where XXX is the name of the inheritable extensible attribute. Note that the column **EA-XXX** displays the name of the inheritable extensible attribute and its value whereas **EAIherited-XXX** displays the inheritance state, which is either **Inherit** or **Override**. Extensible attributes with the following inheritance states will be exported: **Inherited**, **Native**, and **Overridden**.

Note the following about inheritable extensible attributes:

- By default, the value is displayed as **Override** for inheritable extensible attributes, which are at the top of the inheritance chain or if the value of the inherited extensible attribute is overridden at the descendant level.
- If the value is inherited by the descendants of the parent object, then the inheritance state is set to **Inherit**.
- If an extensible attribute is not inheritable or if the status is set to **Not Inherited**, then the **EAIherited-XXX** column will not be displayed for these attributes.

- NIOS does not support EA inheritance for DNS objects, but you might see the **EAlnherited-XXX** column in the CSV file when you export data through Grid Manager. Note that NIOS treats these objects as normal extensible attributes even when you enable inheritance for extensible attributes.

## Importing Inheritable Extensible Attributes

You can specify new inheritable extensible attributes in the spreadsheet and import this file using the CSV Import feature. When you import inheritable extensible attributes using the CSV Import feature, you must specify two new columns, **EA-XXX** and **EAlnherited-XXX**, with relevant values as mentioned in the following tables. Note that *XXX* is the name of the inheritable extensible attribute.

For example, if you want to update the value of an existing inheritable extensible attribute **Building**, you must set the inheritance state to **OVERRIDE** in the spreadsheet. The following example shows that the original value of attribute **Building**, which is replaced by **Millennium Tower**.

A	B	C	D	E
<b>HEADER-NETWORK</b>	<b>ADDRESS*</b>	<b>NETMASK*</b>	<b>EA-Building</b>	<b>EAlnherited-Building</b>
NETWORK	10.251.133.128	255.255.255.192	Millennium Tower	OVERRIDE

The following example shows two new columns, **EA-Region** and **EAlnherited-Region**, which are added to the spreadsheet to associate a new inheritable extensible attribute with an existing object:

A	B	C		
<b>HEADER-NETWORK</b>	<b>ADDRESS*</b>	<b>NETMASK*</b>	<b>EA-Region</b>	<b>EAlnherited-Region</b>
NETWORK	10.251.133.128	255.255.255.192	San Pablo	INHERIT

Note the following about inheritance states:

- When you import attributes for a parent object, the inheritance state must be set to **OVERRIDE**.
  - For descendants, the inheritance state can be set to **OVERRIDE** or **INHERIT**. If you specify **INHERIT**, the attribute value will be inherited from the parent object. If you specify **OVERRIDE**, the original value of the attribute will be replaced with the value specified in the spreadsheet.
- This is valid for Network related objects only. The supported inheritance chain is: **Network View -> Network Container -> Network -> Range -> Host/Fixed Address/Reservation**.

## Importing Active Directory Domains and Sites

The Infoblox CSV format does not support extensible attributes that contain information about Active Directory domains and sites or objects that represent Active Directory domains or sites. The appliance displays an error message when you define values for such extensible attributes in the imported CSV file.

When you export networks, the appliance does not include extensible attributes that contain information about Active Directory domains or sites in the generated .CSV file.

## CSV Import for Topology Rulesets and Rules in DNS Traffic Control

You can import DNS Traffic Control topology rulesets and their rules using the **CSV Import** feature. You must specify the topology rulesets and rules separated by commas in the CSV file:

```
header-dtctopology,name*,comment
header-dtctopologyrule,dest_link*,dest_type*,name*,parent*,sources*,position
dtctopology,topo_server1,TopologyRule
dtctopologyrule,dtc_s1,Server,Rule1,topo_server1,SUBNET/IS_NOT/
172.31.0.0/16,1 dtctopologyrule,pool12,Server,Rule12,topo_server1,"COUNTRY/IS/
Canada",2
```

## Importing Topology Rulesets and Rules

To import topology rulesets, you must specify the **header-topology**, **name** and **comment** columns in the spreadsheet. The **name** column indicates the name of the topology ruleset. To import topology rules, specify the following in the spreadsheet:

A	B	C				
HEADER-DTCTOPOL OGY	NAME*	COMMENT				
HEADER-DTCTOPOL OGYRULE	DEST_LINK*	DEST_TYPE*	NAME*	PARENT*	SOURCES*	POSITION
DTCTOPOL OGY	topo_server1	Topology Server 1				
DTCTOPOL OGYRULE	dtc_s1	SERVER	RULE1	topo_server1	SUBNET/IS/10.0.0.0/8	1
DTCTOPOL OGYRULE	dtc_s2	SERVER	RULE2	topo_server1	SUBNET/IS/10.120.0.0/16	2
DTCTOPOL OGY	topo_pool1	Topology Pool 3				
DTCTOPOL OGYRULE	pool1	POOL	RULE3	topo_pool1	COUNTRY/IS/Antarctica,CONTINENT/IS/Africa,SUBDIVISION/IS/Aden	1

You must specify the **dest\_link**, **dest\_type**, **name**, **parent**, **sources**, and **position** columns when you import a CSV file with topology rules. Note that the **dest\_link** indicates the name of the destination, which is either a pool or server and **position** indicates the order of rules in a topology ruleset. The values that you specify for **dest\_link** must exist in the database. The **dest\_type** indicates the destination type, which can either be a server or a pool. Specify a name for the topology rule in the **name** column. In the **parent** column, you can specify the name of the DTC topology ruleset. The **sources** column must contain either a subnet IP address or a geographic location. The appliance displays an error message if you do not specify valid GeolIP labels when you import a CSV file.

## Supported Object Types

This section describes the supported object types and their corresponding fields for CSV import and export. It also includes examples of how to create data files. Ensure that you review this information before you import or export a data file.

 **Note**

All inherited fields follow the override conventions described in [Data Specific Guidelines](#).

## Supported DNS Object Types

DNS Object Type	Required Fields & Syntax	IDN Supported (Yes/No)
Grid DNS Objects	<a href="#">Grid DNS Objects</a>	Yes
Member DNS Objects	<a href="#">Member DNS Objects</a>	Yes
Authoritative Zone	<a href="#">Authoritative Zone</a>	Yes
Forward-Mapping Zone	<a href="#">Forward-Mapping Zone</a>	Yes
Stub Zone	<a href="#">Stub Zone</a>	Yes
Delegated Zone	<a href="#">Delegated Zone</a>	Yes
Authoritative Name Server Group	<a href="#">Authoritative Name Server Group</a>	Yes
Forwarding Member Name Server Group	<a href="#">Forwarding Member Name Server Group</a>	Yes
Stub Member Name Server Group	<a href="#">Stub Member Name Server Group</a>	Yes
Forward/Stub Server Name Server Group	<a href="#">Forward/Stub Server Name Server Group</a>	Yes
A Record	<a href="#">A Record</a>	Yes
AAAA Record	<a href="#">AAAA Record</a>	Yes
CNAME Record	<a href="#">CNAME Record</a>	Yes
DNAME Record	<a href="#">DNAME Record</a>	Yes
MX Record	<a href="#">MX Record</a>	Yes
NAPTR Record	<a href="#">NAPTR Record</a>	Yes
NS Record	<a href="#">NS Record</a>	Yes (supports only FQDN)
PTR Record	<a href="#">PTR Record</a>	Yes

DNS Object Type	Required Fields & Syntax	IDN Supported (Yes/No)
SRV Record	<a href="#">SRV Record</a>	Yes
TXT Record	<a href="#">TXT Record</a>	Yes
TLSA Record	<a href="#">TLSA Record</a>	No
CAA Record	<a href="#">CAA Record</a>	Yes
Host Record	<a href="#">Host Record</a>	Yes
IPv4 Host Address	<a href="#">IPv4 Host Address</a>	No
IPv6 Host Address	<a href="#">IPv6 Host Address</a>	No
Bulk Host	<a href="#">Bulk Host</a>	No
Rulesets	<a href="#">NXDOMAIN and Blacklist Rulesets</a>	No
NXDOMAIN Rule	<a href="#">NXDOMAIN Rule</a>	No
Blacklist Rule	<a href="#">Blacklist Rule</a>	No
Whitelist Rule	<a href="#">Whitelist Rule</a>	No
DNS64 Synthesis Group	<a href="#">DNS64 Synthesis Group</a>	No
Response Policy Zone	<a href="#">Response Policy Zone</a>	No
Response Policy A Record	<a href="#">Response Policy A Record</a>	No
Response Policy AAAA Record	<a href="#">Response Policy AAAA Record</a>	No
Response Policy IP A Record	<a href="#">Response Policy IP A Record</a>	No
Response Policy IP AAAA Record	<a href="#">Response Policy IP AAAA Record</a>	No
Response Policy MX Record	<a href="#">Response Policy MX Record</a>	No
Response Policy NAPTR Record	<a href="#">Response Policy NAPTR Record</a>	No
Response Policy PTR Record	<a href="#">Response Policy PTR Record</a>	No

DNS Object Type	Required Fields & Syntax	IDN Supported (Yes/No)
Response Policy SRV Record	<a href="#">Response Policy SRV Record</a>	No
Response Policy TXT Record	<a href="#">Response Policy TXT Record</a>	No
Response Policy CNAME Record	<a href="#">Response Policy CNAME Record</a>	No
Response Policy IP Address	<a href="#">Response Policy IP Address</a>	No
Response Policy Client IP Address	<a href="#">Response Policy Client IP Address</a>	No
Response Policy IP Address CNAME	<a href="#">Response Policy IP Address CNAME</a>	No
Response Policy Client IP Address CNAME	<a href="#">Response Policy Client IP Address CNAME</a>	No
Dynamic Update Group	<a href="#">Dynamic Update Group</a>	No
Dynamic Update Cluster Group	<a href="#">Dynamic Update Cluster Group</a>	No

## Supported DHCP Object Types



### Note

IDN is not supported for DHCP object types.

DHCP Object Type	Required Fields & Syntax
Grid DHCP	<a href="#">Grid DHCP Objects</a>
Member DHCP	<a href="#">Member DHCP Objects</a>
Network View	<a href="#">Network View</a>
DNS View	<a href="#">DNS View</a>
IPv4 Network Container	<a href="#">IPv4 Network Container</a>
IPv4 Network	<a href="#">IPv4 Network</a>
IPv6 Network Container	<a href="#">IPv6 Network Container</a>
IPv6 Network	<a href="#">IPv6 Network</a>

<b>DHCP Object Type</b>	<b>Required Fields &amp; Syntax</b>
IPv4 Shared Network	<i>IPv4 Shared Network</i>
IPv6 Shared Network	<i>IPv6 Shared Network</i>
IPv4 DHCP Range	<i>IPv4 DHCP Range</i>
IPv6 DHCP Range	<i>IPv6 DHCP Range</i>
IPv4 Fixed Address and Reservation	<i>IPv4 Fixed Address/Reservation</i>
IPv6 Fixed Address	<i>IPv6 Fixed Address</i>
DHCP Fingerprint	<i>DHCP Fingerprint</i>
DHCP MAC Filter	<i>DHCP MAC Filter</i>
MAC Filter Address Item	<i>MAC Filter Address</i>
Option Filter	<i>Option Filter</i>
Option Filter Match Rule	<i>Option Filter Match Rule</i>
DHCP Fingerprint Filter	<i>DHCP Fingerprint Filter</i>
Relay Agent Filter	<i>Relay Agent Filter</i>
NAC Filter	<i>DHCP Fingerprint Filter</i>
IPv4 Option Space	<i>IPv4 Option Space</i>
IPv6 Option Space	<i>IPv6 Option Space</i>
IPv4 Option Definition	<i>IPv4 Option Definition</i>
IPv6 Option Definition	<i>IPv6 Option Definition</i>
Permissions for DNS resources with associated IP addresses in networks and ranges	<i>Permissions for DNS Resources with Associated IP Addresses in Networks and Ranges</i>
DHCP Failover Association	<i>DHCP Failover Association</i>

## Other Supported Objects

Other Supported Objects	Required Fields and Syntax
Grid Member	For more information on Grid member section, see <a href="#">DHCP Failover Association</a> .
Upgrade Groups Distribution Schedules Upgrade Schedules	<a href="#">Upgrade Groups and Schedules</a>
Named ACLs (access control lists)	<a href="#">Named ACLs</a> <a href="#">ACES in Named ACLs</a>
Infoblox Network Insight	<a href="#">Discovery Credentials</a>

## Grid DNS Objects

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-GridDNS	String	Yes			Identifies the first row as a header row for the Grid DNS objects. Example: GridDNS
refresh	Unsigned integer	No			Indicates the refresh time in seconds. Example: 10800
retry	Unsigned integer	No			Indicates the retry time in seconds. Example: 3600
expire	Unsigned integer	No			Indicates the expiration time in seconds. Example: 2419200
default_ttl	Unsigned integer	No			Indicates the default TTL value in seconds. Example: 28800
negative_ttl	Unsigned integer	No			Indicates the negative TTL value in seconds. Example: 900
lame_ttl	Unsigned integer	No			Indicates the lame TTL value in seconds. Example: 600
email	String	No			Indicates the email address. Example: admin@xyz.com
enable_secondary_notify	Boolean	No			Enable Grid secondaries to send notification. Example: False
enable_notify_source_port	Boolean	No			Enable notification source port Example: False



Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
notify_source_port	Unsigned integer	No			Indicates notify-source port number.
enable_query_source_port	Boolean	No			Enable query source port. Example: False
query_source_port	Unsigned integer	No			Indicates query-source port number.
allow_transfer	ACL	No	Allow zone transfers to	allow_transfer	List of <b>address_tsig_ac</b> items. Example: NACL1 or "12.0.0.12/Deny,1234::/64/Allow". Note that you can import the name of a named ACL in this field.
excluded_servers	IP address list	No			List of excluded servers for zone transfers.
zone_transfer_format_option	String	No			Indicates the zone transfer format. Example: MANY_ANSWERS
allow_query	ACL	No	Allow queries from	allow_query	List of <b>address_tsig_ac</b> items. It can be an IP address, a network entry, Any or a TSIG-/permission. If the first value is not <b>Any</b> or <b>TSIG-</b> , it is assumed to be an IP address or a network entry. Example: 10.0.0.10/Allow, 11.0.0.0/16/Deny, TSIG-foo/xyz/Allow. It can also be a named ACL. Example: NACL1.
recursion_enabled	Boolean	No			Indicates the flag to respond to recursive queries. Example: False
recursive_query_list	ACL	No			It can be an IP address, a network entry, Any or a TSIG-/permission. If the first value is not <b>Any</b> or <b>TSIG-</b> , it is assumed to be an IP address or a network entry. Example: 10.0.0.10/Allow, 11.0.0.0/16/Deny,TSIG-foo/xyz/Allow. It can also be a named ACL. Example: NACL1.
allow_update	ACL	No	Allow updates from	allow_update	List of <b>address_tsig_ac</b> items. It can be an IP address, a network entry, Any or a TSIG-/permission. If the first value is not <b>Any</b> or <b>TSIG-</b> , it is assumed to be an IP address or a network entry. Example: 10.0.0.10/Allow, 11.0.0.0/16/Deny, TSIG-foo/xyz/Allow. It can also be a named ACL. Example: NACL1.
allow_update_forwarding	Boolean	No	Allow updates from	forward_to	Enable update forwarding for secondary zones. Example: False
allow_bulkhost_ddns	String	No			Enable updates to PTR records sourced from a bulkhost. Example: Refuse

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
forwarders_only	Boolean	No	Use Forwarders Only		Enable use of forwarders only. Example: False
allow_forwarder	IP address list	No			Indicates the list of forwarders.
enable_custom_root_server	Boolean	No			Indicates the flag to enable custom root servers. Example: False
root_name_servers	Root nameserver list	No			Indicates the list of custom root servers. Example: mm1.test.com/1.1.1.1/,... The appliance displays an error message if the <b>root_name_servers</b> column has an empty value when the <b>enable_custom_root_server</b> field is set to <b>True</b> in the imported CSV file.
enable_blackhole	Boolean	No			Enable blackhole setting. Example: False
blackhole	ACL	No			Indicates the list of banned addresses. Example: "NACL" or "12.0.0.12/Deny,1234::/64/Allow,.."
notify_delay	Unsigned integer	No		notify_delay	This field specifies the seconds of delay the notify messages are sent to the secondaries. The valid value is between 5 and 86400 seconds. Example: 5
enable_nxdomain_redirect	Boolean	No			Enable intercept and redirect nxdomain responses. Example: False
nxdomain_redirect_addresses	IP address list	No			Indicates the list of IPv4 addresses to redirect to for nxdomain responses. Example: "1.1.1.1,2.2.2.2,..."
nxdomain_redirect_ttl	Unsigned integer	No			Indicates the NXDOMAIN redirect ttl in seconds. Example: 60
nxdomain_log_query	Boolean	No			If you set this to <b>True</b> , the appliance logs the NXDOMAIN redirections. Example: False
nxdomain_rulesets	Pattern list	No			Indicates the list of ruleset objects that are used for NXDOMAIN redirection. Example: pattern1/MODIFY, pattern2/PASS, ...
enable_blacklist	Boolean	No		enable_blacklist	Enable or disable blacklist redirection at the Grid level. Example: False

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
blacklist_redirect_addresses	IP address list	No		blacklist_redirect_addresses	Indicates the list of IPv4 addresses to which the blacklisted queries are redirected. Example: 1.1.1.1,2.2.2.2
blacklist_action	String	No	Action	blacklist_action	Indicates the action to be performed when a domain name matches the pattern defined in an assigned rule. Example: Refuse
blacklist_redirect_ttl	Unsigned integer	No		blacklist_redirect_ttl	Indicates the blacklist redirect TTL value in seconds. Example: 60
blacklist_log_query	Boolean	No		blacklist_log_query	When this is set to <b>True</b> , blacklisted queries are logged. Example: False
blacklist_rulesets	List of domain names	No		blacklist_rulesets	List of ruleset objects that are used for blacklist redirection. Example: list1.com, list2.com, ...
enable_dns64	Boolean	No			Enable DNS64 synthesis. Example: False
dns64_groups	List of Dns64 groups	No			List of SynthesisGroup objects. Example: dns64_groupA, dns64_groupB, ...
host_rrset_order	Boolean	No			Specify <b>True</b> to set the <b>enable_host_rrset_order</b> flag or <b>False</b> to deactivate <b>enable_host_rrset_order</b> value at the Grid level. Example: False
preserve_host_rrset_order_on_secondaries	Boolean	No			Specify <b>True</b> to set the <b>preserve_host_rrset_order_on_secondaries</b> flag or <b>False</b> to deactivate <b>preserve_host_rrset_order_on_secondaries</b> value at the Grid level. The default value is <b>False</b> . Example: False
filter_aaaa	String	No			Indicates the type of AAAA filtering for this Grid DNS object. The default value is <b>No</b> . Example: Yes
filter_aaaa_list	ACL	No			Indicates the list of IPv4 addresses and networks from which queries are received. Note that the AAAA filtering is applied to these addresses. Example: "12.0.0.12/Deny,13.0.0.0/8/Allow..." or "NACL1"
copy_xfer_to_notify	Boolean	No			Enable or disable copying of the allowed IP addresses from zone transfer list into also-notify statement in named.conf. Example: False

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
transfers_in	Unsigned integer	No			Indicates the number of maximum concurrent transfers for the Grid. You can specify unsigned integers between 10 and 100. The default value is 10. Example: 10
transfers_out	Unsigned integer	No			Indicates the number of maximum outbound concurrent zone transfers for the Grid. You can specify unsigned integers between 10 and 100. The default value is 10. Example: 10
transfers_per_ns	Unsigned integer	No			Indicates the number of maximum concurrent transfers per member for the Grid. You can specify unsigned integers between two and 100. The default value is two. Example: 2
serial_query_rate	Unsigned integer	No			Indicates the number of maximum concurrent SOA queries per second for the Grid. You can specify unsigned integers between 20 and 100. The default value is 20. Example: 20
max_cache_ttl	Unsigned integer	No			Indicates the maximum time (in seconds) for which the server will cache positive answers. The default value is 604800.
max_ncache_ttl	Unsigned integer	No			Indicates the maximum time (in seconds) for which the server will cache negative (NXDOMAIN) responses. The default value is 10800. The maximum allowed value is 604800.
disable_edns	Boolean	No			Enable or disable EDNS0 support for queries that require recursive resolution. The default value is <b>False</b> .
query_rewrite_enabled	Boolean	No			When this is set to <b>True</b> , query rewrite is enabled at the Grid level. Example: False
query_rewrite_domain_names	List of domain names	No			Indicates the list of domain names that trigger DNS query rewrite. Example: "aa.com, bb.com."
query_rewrite_prefix	String	No			Indicates the domain name prefix for DNS query rewrite. The default value is <b>undefined</b> .
rpz_drop_ip_rule_enabled	Boolean	No	Ignore RPZ-IP triggers with too small prefix lengths		When this is set to <b>True</b> , DNS server ignores RPZ-IP rules with prefix lengths that are less than the specified prefix length limit. Example: TRUE

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
rpz_drop_ip_rule_min_prefix_length_ipv4	Unsigned Integer	No	Minimum IPv4 Prefix Length		Indicates the minimum IPv4 prefix length for RPZ-IP triggers. The default value is 29.
rpz_drop_ip_rule_min_prefix_length_ipv6	Unsigned Integer	No	Minimum IPv6 Prefix Length		Indicates the minimum IPv6 prefix length for RPZ-IP triggers. The default value is 112.
rpz_hit_rate_interval	Unsigned Integer	No	Interval		Indicates the minimum time interval in seconds between RPZ hit rate checks. The default interval is 10 seconds.
rpz_hit_rate_min_query	Unsigned Integer	No	Minimum query		Indicates the minimum number of queries between RPZ hit rate checks. The default value is 1000.
rpz_hit_rate_max_query	Unsigned Integer	No	Maximum query		Indicates the maximum number of queries between RPZ hit rate checks. The default value is 100000.

## Authoritative Zone

You can import the name of a named ACL in the fields that support named ACLs, such as `allow_transfer`, `allow_query`, and `allow_update`.



### Note

IDN is supported for object types: `fqdn`, `soa_mname`, and `soa_email`. You can use punycode or IDNs while importing these objects.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-AuthZone	String	Yes			
fqdn	FQDN	Yes	Name	name	Example: test.com
zone_format	String	Yes			Valid values are FORWARD, IPV4, and IPV6
view	String	No			If no view is specified, the Default view is used.
prefix	String	No			Prefix is used for reverse-mapping RFC2317 zones only. If you include a prefix in a forward-mapping zone, the appliance ignores the prefix. No error message is generated.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
_new_prefix	String	No			Add this field to overwrite the prefix field when you select the overwrite or merge option. Use the hostname of the grid member in this field. Example: infoblox.localdomain
is_multimaster	Boolean	No	Multi-master	is_multimaster	Indicates whether the zone has multiple primary servers. Example: True
grid primaries	Grid member list and stealth state	No	Grid Primary/Stealth	primary stealth	Data must be in the following format: "hostname/stealth" Example: "foo.localadmin/False,corp1.com/True,..."
external primaries	Server list	No	External Primary	primary	Data must be in the following format: "name/ip/stealth/use_2x_tsig/ use_tsig/ tsig_name/tsig_key/ tsig_key_algorithm". Only name and ip are required fields. If no value is specified for stealth, use_2x_tsig, and use_tsig, the default value FALSE is used. If either use_2x_tsig or use_tsig is TRUE, tsig_name and tsig_key are required. If no value is specified for tsig_key_algorithm, the default value is HMAC-MD5. If both use_2x_tsig and use_tsig are TRUE, only use_tsig = TRUE and the tsig key name and key are imported. Example: "ext1.test.com/1.1.1.1/FALSE"
grid secondaries	Member server list	No	Grid Secondary	secondaries	Data must be in the following format: "hostname/stealth/lead/grid_replicate". Only hostname is required. If you do not specify values for stealth, lead, and grid_replicate, the default value FALSE is used. Example: "member1.localdomain/FALSE/TRUE/FALSE"

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
external_secondaries	Server list	No	External Secondary	secondaries	Data must be in the following format: "name/ip/stealth/use_2x_tsig/use_tsig/tsig_name/tsig_key/tsig_key_algorithm". Only name and ip are required fields. If no value is specified for stealth, use_2x_tsig, and use_tsig, the default value FALSE is used. If either use_2x_tsig or use_tsig is TRUE, tsig_name and tsig_key are required. If no value is specified for tsig_key_algorithm, the default value is HMAC-MD5. If both use_2x_tsig and use_tsig are TRUE, only use_tsig = TRUE and the tsig key name and key are imported. Example: "sec1.com/1.1.1.1/FALSE/FALSE/FALSE/foo/sdfssdf86ew"
ns_group	String	No	Name server group	ns_group	Authoritative name server group name. Example: name-ns-group1
comment	String	No	Comment	comment	
disabled	Boolean	No	Disable	disable	Example: FALSE
create_underscore_zones	Boolean	No	Automatically create underscore zones	create_underscore_zones	Example: FALSE
allow_active_dir	List of IP addresses	No	Allow unsigned updates from these domain controllers	enable_ad_server	The Valid value is a list of IP addresses. Example: "1.1.1.1, 10.0.0.1"
soa_refresh	Unsigned integer	No	Refresh	soa_refresh	When you modify this field to override an inherited value, you must include values for all SOA timer fields. The appliance updates all the SOA timers when you update any of them.
soa_retry	Unsigned integer	No	Retry	soa_retry	Ensure that you include this field when you override the soa_refresh field.
soa_expire	Unsigned integer	No	Expire	soa_expire	Ensure that you include this field when you override the soa_refresh field.
soa_default_ttl	Unsigned integer	No	Default TTL	soa_default_ttl	Ensure that you include this field when you override the soa_refresh field.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
soa_negative_ttl	Unsigned integer	No	Negative-caching TTL	soa_negative_ttl	Ensure that you include this field when you override the soa_refresh field.
soa_mnames	FQDN list	No	List of SOA MNAME fields	soa_mname	Data must include the FQDN and hostname Example: "foo.localdomain/foobar.localadmin,..."
soa_email	Email address	No	Email address for SOA MNAME field	soa_email	Example: root@
soa_serial_number	Unsigned integer	No	Serial Number	soa_serial_number	
disable_forwarding	Boolean	No	Don't use forwarders...	disable_forwarding	Example: TRUE
allow_update_forwarding	Boolean	No	Allow updates from	forward_to	Example: FALSE
update_forwarding	ACL	No	Allow updates from... Permission table	forward_to	Data must be in the following formats: ip address/permissionnetwork/networkcidr/permissionANY/permissionTSIG-XXX/permission  Permission can be ALLOW or DENY If the first value is not Any or TSIG-, it is assumed to be an IP or network address. Example: "10.0.0.10/Allow,11.0.0.0/16/Deny,TSIG-foo/sdfsfwbsdgsfsw8sdf/Allow"
allow_transfer	ACL	No	Allow zone transfers to	allow_transfer	Example: "12.0.0.12/Deny,1234.:/64/Allow" Note that you can import the name of a named ACL in this field.
allow_update	ACL	No	Allow updates from	allow_update	Example: "13.0.0.0/8/Allow" Note that you can import the name of a named ACL in this field.
allow_query	ACL	No	Allow queries from	allow_query	Example: "127.0.0.1/Allow" Note that you can import the name of a named ACL in this field.



Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
notify_delay	Unsigned integer	No		notify_delay	This field specifies the seconds of delay the notify messages are sent to the secondaries. The valid value is between 5 and 86400 seconds. Example: 10
EA-Site	String	No	Extensible attribute	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	List	No	Extensible attribute	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role		ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for DNS zones. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding DNS Zones

This example shows how to add a forward mapping zone, corp100.com, with a Grid primary and a Grid secondary, where the grid secondary = hostname/stealth/lead/grid\_replicate.

```
header-authzone, fqdn*, zone_format*, comment, grid_secondaries
authzone, corp100.com, FORWARD, USA, member.infoblox.com/False/ False3
```

This example shows how to create a data file to add an IPv4 reverse mapping zone, 100.0.0.0/8, with a Grid primary and an external secondary, where the external secondary = name/ip/stealth/use\_2x\_tsig/use\_tsig/ tsig\_name/tsig\_key.

```
header-authzone, fqdn*, zone_format*, external_secondaries
authzone, 100.0.0.0/8, IPV4, ns2.com/2.2.2.2/False/False/False/None/None
```

This example shows how to create a data file to add an IPv6 reverse-mapping zone, 1234::/64, with an external primary and a Grid secondary, where the external primary = name/ip/stealth/use\_2x\_tsig/use\_tsig/ tsig\_name/tsig\_key.

```
header-authzone, fqdn*, zone_format*, external primaries, grid_secondaries
authzone, 1234::/64, IPV6, ns1.com/1.1.1.1/False/False/False/None/None,
member.infoblox.com/False/False/False
```

### Overwriting DNS Zone Data

This example shows how to overwrite a comment from "USA" to "Japan" and remove the Grid secondary.

```
header-authzone, fqdn*, zone_format*, comment, grid_secondaries
authzone, corp100.com, FORWARD, Japan
```

### Merging DNS Zone Data

This example shows how to merge the extensible attribute "Site" = "HQ" and add the RW permission to an admin group "DNS\_admins".

```
header-authzone, fqdn*, zone_format*, ADMGRP-DNS_admins, EA-site
authzone, corp100.com, FORWARD, RW, HQ
```

This example shows how to add an external secondary with these values: "ns3.com/2.2.2.2/False/False/False/None/None".

```
header-authzone, fqdn*, zone_format*, external_secondaries
authzone, 100.0.0.0/8, IPV4, ns2.com/2.2.2.2/False/False/False/None/None, ns3.com/
2.2.2.2/ False/False/False/None/None
```

### Adding Named ACL Data

This example shows how to import the names of named ACLs in supported fields, such as allow\_transfer, allow-query, and allow\_update:

```
Header-
authzone, fqdn*, grid_primary, view, external_secondaries, allow_transfer, allow_query, zone_type, allow_active_dir, allow_update, zone_format, notify_delay, disabled, grid_primary_stealth, soa_negative_ttl, soa_mname, soa_default_ttl, soa_retry, , create_under_score_zones, soa_serial_number, soa_email, comment, soa_expire, soa_refresh
authzone, test_data_export.com, infoblox.localdomain, default, test_data.infoblox.com/1.1.1.1/TRUE, "12.0.0.12/Deny, 1234::/64/Allow", My_Named_ACL, Authoritative, 1.2.3.4, "1234::/64/Allow", FORWARD, 100, FALSE, FALSE, 100, mname2, 300, 600, FALSE, FALSE, 1, soaemail@infoblox.c (mailto:%2CFORWARD%2C100%2CFALSE%2CFALSE%2C100%2Cmname2%2C300%2C600%2CFALSE%2CFALSE%2C1%2Csoaemail@infoblox.co) o m, Authzone2, 200, 500
```

```

authzone, test_csv_export.com, infoblox.localdomain, default, test_csv.infoblox.com
/1.1.1. 1/TRUE, My_Named_ACL, "12.0.0.12/Deny, 1234::/64/
Allow", , 2.3.4.5, "4321::/64/Allow", FORWARD
D, 100, FALSE, FALSE, 400, mname1, 900, 800, FALSE, FALSE, 1, tel@infoblox.com, Authzone1, 10
0, 200

```

## Delegated Zone

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-DelegatedZone	String	Yes			Identifies the first row as a header row for delegated zones. Example: DelegatedZone
fqdn	FQDN	Yes	Name	zone	This field combines the AAAA record name and the zone name to form the FQDN. Example: aaa1.corp100.com
view	String	No	DNS View	views	If no view is specified, the default view is used. Example: Default
zone_format	String	Yes	Type		Valid values are <b>FORWARD</b> , <b>IPV4</b> , and <b>IPV6</b> .
prefix	String	No	RFC 2317 Prefix	prefix	Prefix is used for reverse-mapping RFC2317 zones only. If you include a prefix in a forward-mapping zone, the appliance ignores the prefix. No error message is generated.
disabled	Boolean	No	Disable	disable	Enable or disable the zone. Example: FALSE
comment	String	No	Comment	comment	Example: Delegated zone header.
delegate_to	Delegated Servers list	Yes			Example: delegate_server1.test.com/1.1.1.1/,
delegated_ttl	Unsigned integer	No			This is an inherited field. Example: 28800
EA-Site	String	No	Extensible attribute	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. Example: California.
EA-Users	List	No	Extensible attribute	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. Example: ['Annie', 'John'].
ADMGRP-XXXX	String	No	Permissions Admin Group/Role	permission	ADMGRP-JimSmith is an example of an admin permission of a specific admin group. Example: RW

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ns_group	String	No	Name server group	ns_group	Authoritative name server group name. Example: name-ns-group1
_new_prefix	String	No			Add this field to overwrite the prefix field when you select the Overwrite or Merge option. Use the host name of the Grid member in this field. Example: infoblox.localdomain
ddns_protected	Boolean	No	Protected		Add this field in order to restrict DDNS updates to record.
ddns_principal	String	No	Principal		Displays the principal name for dynamic records with the GSS-TSIG principal authentication enabled in the advanced updates properties of the Grid, view, zone, or Standalone.

## Delegation Name Server Group

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-delegationnsgroup	String	Yes			Identifies the first row as a header row for the delegation name server group objects. Example: DelegationNsGroup.
group_name	String	Yes	Name		Indicates the name of the delegation name server group. Example: ns_group1
_new_group_name	String	No			You can overwrite the group name.
delegate_to	String/IP Address	Yes	Name Server/ Address		List of name servers with valid IP address. Example: "foo.com/1.1.1.1"
comment	String	No	Comment	Comment	Example: This is a delegation name server group.

## Forwarding Member Name Server Group

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-ForwardingMemberNsGroup	String	Yes			Identifies the first row as a header row for the forwarding member name server group objects. Example: ForwardingMemberNsGroup.
group_name	String	Yes	Name	name	Indicates the name of the forwarding member name server group. Example: fwd_ns_group1
_new_group_name	String	No			You can overwrite the group name.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
comment	String	No	Comment	comment	Example: This is a forwarding member name server group.
forwarding_servers	Forwarding members list	Yes	Edit-Per-Member Forwarders Editor	forwarding_servers	List of forwarding servers per Grid member. Example: False/True/infoblox.localdomain/[test/2.2.2.2], where: - False indicates the <b>Use Forwarders Only</b> checkbox is not selected. - True indicates the <b>Override Default Forwarders</b> checkbox is selected. - test/2.2.2.2 - Custom forwarders <b>Note:</b> You cannot clear the custom forwarders using the CSV import operation.
EA-XXX	String	No	Extensible attribute	extensible_attributes	EA-XXX is an example of a user defined attribute. You can add other user defined attributes to the data file. Example: John.

## Forward or Stub Server Name Server Group

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-ForwardStub ServerNsGroup	String	Yes			Identifies the first row as a header row for the forward/stub server name server group objects. Example: ForwardStubServerNsGroup.
group_name	String	Yes	Name	name	Indicates the name of the forward/stub server name server group. Example: ext_ns_group1
_new_group_name	String	No			You can overwrite the group name.
comment	String	No	Comment	comment	Example: This is a forward/stub server name server group.
external_servers	External server list	Yes		external_servers	List of external servers.
EA-XXX	String	No	Extensible attribute	extensible_attributes	EA-XXX is an example of a user defined attribute. You can ad

## Stub Member Name Server Group

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-StubMemberNsGroup	String	Yes			Identifies the first row as a header row for the stub member name server group objects. Example: StubMemberNsGroup.
group_name	String	Yes	Name	name	Indicates the name of the stub member name server group. Example: stub_ns_group1
_new_group_name	String	No			You can overwrite the group name.
comment	String	No	Comment	comment	Example: This is a stub member name server group.
stub_members	Member server list	Yes		stub_members	List of stub Grid members.
EA-XXX	String	No	Extensible attribute	extensible_attributes	EA-XXX is an example of a user defined attribute. You can add other user defined attributes to the data file. Example: John.

## A Record



### Note

IDN is supported for object type: fqdn. You can use IDN or punycode while importing this object.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-ARecord	String	Yes			Example: ARecord
fqdn	FQDN	Yes	Name	name	This field combines the A record name and the zone name to form the FQDN. Example: a1.corp100.com
_new_fqdn	FQDN	No			Add this field to overwrite the fqdn field when you select the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used.
address	IP address	Yes	IP Address	ipv4addrss	Example: 192.138.1.1
_new_address	IP address	No			Add this field to overwrite the address field when you select the overwrite or merge option.
comment	String	No	Comment	comment	

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
disabled	Boolean	No	Disable	disable	Example: FALSE
ttd	Unsigned integer	No	TTL	ttd	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 7200
EA-Site	String	No	Extensible attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	List	No	Extensible attribute Users	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for A records. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding an A Record

This example shows how to add an A record, bind\_a.corp100.com, with the extensible attribute Site = Infoblox, and the permission, DNS\_Admins = RO.

```
header-arecord, address*, fqdn*, ADMGRP-DNS_Admins, EA-Site
```

```
arecord, 100.0.0.1, bind_a.corp100.com, RO, Infoblox
```

### Overwriting A Record Data

This example shows how to modify the permission of the admin group DNS\_Admins from RO to DENY in an existing A record, bind\_a.corp100.com.

```
header-arecord, address*, fqdn*, ADMGRP-DNS_Admins
```

```
arecord, 100.0.0.1, bind_a.corp100.com, DENY
```

### Merging DNS Zone Data

This example shows how to merge the TTL value = 1280 to an existing A record, bind\_a.corp100.com.

```
header-arecord, address*, fqdn*, ttl
```

```
arecord, 100.0.0.1, bind_a.corp100.com, 1280
```

## AAAA Record

 **Note**

IDN is supported for object type: `fqdn`. You can use IDN or punycode while importing this object.

Field Name	Data Type	Required (Yes/No)	Associated GUI field	Associated PAPI Method	Usage and Guidelines
Header-AaaaRecord	String	Yes			Example: AaaaRecord
fqdn	FQDN	Yes	Name	zone	This field combines the AAAA record name and the zone name to form the FQDN. Example: aaaa1.corp100.com
_new_fqdn	FQDN	No			Add this field to overwrite the fqdn field when you select the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default
address	IPv6 address	Yes	IP Address	ipv6addrss	Example: 100::10
_new_address	IPv6 address	No			Add this field to overwrite the address field when you select the overwrite or merge option.
comment	String	No	Comment	comment	
disabled	Boolean	No	Disable	disabled	Example: FALSE
ttl	Unsigned integer	No	TTL	ttl	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 7200
EA-Site	String	No	Extensible attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	List	No	Extensible attribute Users	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .



Field Name	Data Type	Required (Yes/No)	Associated GUI field	Associated PAPI Method	Usage and Guidelines
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for AAAA records. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding an AAAA Record

This example shows how to add an AAAA record, `bind_aaaa.corp100.com`, with a comment = add by superuser, and TTL = 3600.

```
header-aaaarecord,address*,fqdn*,comment,ttl
aaaarecord,1234:1234::1,bind_aaaa.corp100.co,add by superuser,3600
```

### Overwriting AAAA Record Data

This example shows how to modify an existing AAAA record from address `1234:1234::1` to `1234:1234::2`, and TTL from 3600 to 3800.

```
header-aaaarecord,address*,_new_address,fqdn*,ttl
aaaarecord,1234:1234::1,1234:1234:2,bind_aaaa.corp100.com,3800
```

### Merging AAAA Record Data

This example shows how to disable an existing AAAA record.

```
header-aaaarecord,address*,fqdn*,disabled
aaaarecord,1234:1234::2,bind_aaaa.corp100.com,TRUE
```

## Alias Records



### Note

IDN is supported for object type: `fqdn`. You can use IDN or punycode while importing this object.

Field Name	Data Type	Required (Yes/No)	Associated GUI field	Associated PAPI Method	Usage and Guidelines
Header-aliasrecord	String	Yes	NA	NA	Example: Alias Record

Field Name	Data Type	Required (Yes/No)	Associated GUI field	Associated PAPI Method	Usage and Guidelines
fqdn	FQDN	Yes	Name	name	This field combines the Alias record name and the zone name to form the FQDN. Example: aaaa1.corp100.com
new_fqdn	FQDN	No	NA	dns_name	
target_name	String	Yes	Target Name	target_name	You can type the domain name for the resource. Examples: <ul style="list-style-type: none"> <li>• CloudFront distribution domain name: d111111abcdef8.cloudfront.net</li> <li>• Elastic Beanstalk environment CNAME: example.elasticbeanstalk.com</li> <li>• ELB load balancer DNS name: example-1.us-east-1.elb.amazonaws.com</li> <li>• S3 website endpoint: s3-website.us-east-2.amazonaws.com</li> <li>• Resource record set in this hosted zone: www.example.com</li> </ul>
new_target_name	String	No	NA	dns_target_name	
target_type	String	Yes	Target Type	target_type	Type of the aliased resource record, can be one of: A, AAAA, MX, NAPTR, PTR, SPF, SRV, TXT.
_new_target_type	String	No	NA	NA	
comment	String	No	Comment	comment	Additional Information
ttl	Boolean	No	TTL	ttl	This is an inherited field. For information, see <i>Data Specific Guidelines</i> . Example: 7200
disabled	Boolean	No	Disable	disabled	Example: FALSE
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default
Creator	String	No	NA	creator	

## CNAME Record



### Note

IDN is supported for object types: fqdn and canonical\_name. You can use punycode or IDNs while importing these objects.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-CnameRecord	String	Yes			Example: CnameRecord
fqdn	FQDN	Yes	Alias	zone	This field combines the CNAME record name and the zone name to form the FQDN. Example: c1.corp100.com
_new_fqdn	FQDN	No			Add this field to overwrite the fqdn field when you select the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default
canonical_name	Domain name	Yes	Canonical Name	canonical	Example: www.corp100.com
comment	String	No	Comment	comment	
disabled	Boolean	No	Disable	disabled	Example: FALSE
ttl	Unsigned integer	No	TTL	ttl	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 28800
EA-Site	String	No	Extensible attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for CNAME records. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a CNAME Record

This example shows how to add a CNAME record, bind\_cname.corp100.com, to the Default DNS view.

```
header-cnamerecord, fqdn*, canonical_name*, view
cnamerecord, bind_cname.corp100.com, somewhere.corp100.com, default
```

### Overwriting CNAME Record Data

This example shows how to override a canonical name from somewhere.corp100.com to somewhere2.corp100.com.

```
header-cnamerecord, fqdn*, canonical_name*
cnamerecord, bind_cname.corp100.com, somewhere2.corp100.com
```

## Merging CNAME Record Data

This example shows how to merge the following data: admin group DNS\_Admins with RW permission and extensible attribute Site = New York.

```
header-cnamerecord, fqdn*, ADMGRPDNS_Admins, EA-Site
cnamerecord, bind_cname.corp100.com, RW, New York
```

## DNAME Record



### Note

IDN is supported for object types: fqdn and target. You can use punycode or IDNs while importing these objects.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-DnameRecord	String	Yes			Example: DnameRecord
fqdn	FQDN	Yes	Alias	zone	This field combines the DNAME record name and the zone name to form the FQDN. Example: d1.corp100.com
_new_fqdn	FQDN	No			Add this field to overwrite the fqdn field when you select the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default
target	Domain name	Yes	Target	target	Example: d1.foo.com
comment	String	No	Comment	comment	
disabled	Boolean	No	Disable	disabled	Example: FALSE
ttd	Unsigned integer	No	TTL	ttd	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 28800
EA-Site	String	No	Extensible attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for DNAME records. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a DNAME Record

This example shows how to add a DNAME record, `bind_dname.corp100.com`, with `target = bind_dname.corp200.com`, extensible attribute `Site = HQ`, and `disabled = TRUE`.

```
header-dnamerecord,fqdn*,target*,disabled,EA-Site
dnamerecord,bind_dname.corp100.com,bind_dname.corp200.com,TRUE,HQ
```

### Overwriting DNAME Record Data

This example shows how to enable the existing DNAME record, `bind_dname.corp100.com`, and add admin group `DNS_Admins` with RO permission.

```
header-dnamerecord,fqdn*,target*,disabled,ADMGRPDNS_Admins
dnamerecord,bind_dname.corp100.com,bind_dname.corp200.com,FALSE,RO
```

### Merging DNAME Record Data

This example shows how to add a comment to the existing DNAME record, `bind_dname.corp100.com`, and change the TTL to 3860.

```
header-dnamerecord,fqdn,target*,comment,ttl
dnamerecord,bind_dname.corp100.com,bind_dname.corp200.com,Add by DNS admin,3860
```

## MX Record



### Note

IDN is supported for object types: `fqdn` and `mx`. You can use punycode or IDNs while importing these objects.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-MxRecord	String	Yes			Example: MxRecord
fqdn	FQDN	Yes	Mail Destination	zone	This field combines the MX record name and the zone name to form the FQDN. Example: MX1.corp100.com
_new_fqdn	FQDN	No			Add this field to overwrite the fqdn field when you select the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
mx	Domain name	Yes	Mail Exchange	exchanger	Example: mailer.foo.com
_new_mx	Domain name	No			Add this field to overwrite the mx field when you select the overwrite or merge option.
priority	Unsigned integer	Yes	Preference	pref	Example: 10
_new_priority	Unsigned integer	No			Add this field to overwrite the priority field when you select the overwrite or merge option.
comment	String	No	Comment	comment	
disabled	Boolean	No	Disable	disable	Example: FALSE
ttl	Unsigned integer	No	TTL	ttl	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 28800
EA-Site	String	No	Extensible attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Extensible attribute Country	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for MX records. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding an MX Record

This example shows how to add an MX record, bind\_mx.corp100.com, with a mail exchanger, exchange.corp100.com and priority = 20.

```
header-mxrecord,fqdn*,mx*,priority*
mxrecord,bind_mx.corp100.com,exchange.corp100.com,20
```

## Overwriting MX Record Data

This example shows how to overwrite an existing MX record with a new fqdn, bind\_mx2.corp100.com, and a new mail exchanger, new\_exchange.corp100.com.

```
header-mxrecord, fqdn*, _new_fqdn, mx*, _new_mx
```

```
mxrecord, bind_mx.corp100.com, bind_mx2.corp100.com, exchange.corp100.com, new_exchange.corp100.com
```

## Merging MX Record Data

This example shows how to merge data to the existing MX record, bind\_mx2.corp100.com, by adding inherited TTL value and extensible attributes Site = USA.

```
header-mxrecord, fqdn*, mx*, priority*, ttl, EA-Site
```

```
mxrecord, bind_mx2.corp100.com, new_exchange.corp100.com, 20, USA
```

## NAPTR Record



### Note

IDN is supported for object types: fqdn and replacement. You can use punycode or IDNs while importing these objects.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-NaptrRecord	String	Yes			Example: NaptrRecord
fqdn	FQDN	Yes	Domain	name	This field combines the domain name and the zone name to form the FQDN. Example: aptr1.corp100.com
_new_fqdn	FQDN	No			Add this field to overwrite the fqdn field when you select the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default
order	Unsigned integer	Yes	Order	order	Example: 10
_new_order	Unsigned integer	No			Add this field to overwrite the order field when you select the overwrite or merge option.
preference	Unsigned integer	Yes	Preference	preference	Example: 20

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
_new_preference	Unsigned integer	No			Add this field to overwrite the preference field when you select the overwrite or merge option.
flags	String	No	Flags	flags	You can leave this field empty. Example: U
_new_flags	String	No			Add this field to overwrite the flags field when you select the overwrite or merge option.
services	String	No	Service	services	You can leave this field empty. Example: SIP+D2U
_new_services	String	No			Add this field to overwrite the services field when you select the overwrite or merge option.
regexp	String	No	REGEX	regexp	You can leave this field empty. Example: http://([^\:;]+)\!1!
_new_regexp	String	No			Add this field to overwrite the regexp field when you select the overwrite or merge option.
replacement	String	Yes	Replacement	replacement	Example: corp100.com
_new_replacement	String	No			Add this field to overwrite the replacement field when you select the overwrite or merge option.
comment	String	No	Comment	comment	
disabled	Boolean	No	Disable	disable	Example: FALSE
ttl	Unsigned integer	No	TTL	ttl	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 28800
EA-Site	String	No	Extensible attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Extensible attribute Country	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .



## Examples

This section contains examples of how to create data files for NAPTR records. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a NAPTR Record

This example shows how to add a NAPTR record, `bind_naptr.corp100.com`, with `order = 10`, `preference = 20`, and `replacement = corp200.com`.

```
header-naptrrecord, fqdn*, order*, preference*, replacement*
naptrrecord, bind_naptr.corp100.com, 10, 20, corp200.com
```

### Overwriting NAPTR Record Data

This example shows how to overwrite the FQDN of an existing NAPTR record from `bind_naptr.corp100.com` to `bind_naptr2.corp100.com`.

```
header-naptrrecord, fqdn*, _new_fqdn order*, preference*, replacement*
naptrrecord, bind_naptr.corp100.com, bind_naptr2.corp100.com, 10, 20, corp200.com
```

This example shows how to override the preference of an existing NAPTR record from 20 to 25 and replacement from `corp200.com` to `corp300.com`.

```
header-
naptrrecord, fqdn*, order*, preference*, _new_preference, replacement*, _new_replacem
ent
naptrrecord, bind_naptr.corp100.com, 10 20 25, corp200.com, corp300.com
```

### Merging NAPTR Record Data

This example shows how to merge `Service = http+E2U` and `TTL = 3600 seconds` to an existing NAPTR record.

```
header-naptrrecord, fqdn*, order*, preference*, replacement*, services, ttl
naptrrecord, bind_naptr.corp100.com, 10, 25, Corp300.com, http+E2U, 3600
```

## NS Record

### Note

IDN is supported for object type: fqdn. You can use IDN or punycode while importing this object.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-NsRecord	String	Yes			Example: NsRecord
fqdn	FQDN	Yes	Zone	name	This field combines the domain name and the zone name to form the FQDN. Example: test.corp100.com

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default
dname	Domain name	Yes	Name Server	nameserver	Example: ns1.corp100.com
_new_dname	Domain name	No			Add this field to overwrite the dname field when you select the overwrite or merge option.
zone_nameservers	Server list	Yes	Name server list	addresses	Data must be in the following format: "IPAddress1/auto_create_ptr1, IPAddress2/auto_create_ptr2" Example: "10.0.0.4/TRUE, 10.0.0.44/FALSE, 10.0.0.55/TRUE"



#### Note

When you perform a CSV export of automatically created NS records using Infoblox API, the **zone\_nameservers** field will have an empty value. Therefore, if you import the previously exported CSV file that includes automatically created NS records through the Infoblox GUI, then the CSV import fails and Grid Manager displays an error message.

## Examples

This section contains examples of how to create data files for NS records. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding an NS Record

This example shows how to add an NS record corp100.com in the Default DNS view with DNAME (name server) = ns1.corp100.com, name server address = 100.0.0.101, and TRUE for adding a PTR Record.

```
header-nsrecord, fqdn*, view, dname*, zone_nameservers*
nsrecord, corp100.com, default, ns1.corp100.com, "100.0.0.101/TRUE"
```

This example shows how to add an NS record corp200.com in the Internal DNS view with DNAME (name server) = ns1.corp200.com and two name server addresses: 200.0.0.101 with TRUE for adding a PTR Record and 200.0.0.102 with TRUE for adding a PTR record.

```
header-nsrecord, fqdn*, view, dname*, zone_nameservers*
nsrecord, corp200.com, Internal, ns1.corp200.com, "200.0.0.101/TRUE, 200.0.0.102/TRUE"
```

### Overwriting NS Record Data

This example shows how to overwrite the DNAME of an existing NS record from ns1.corp100.com to ns2.corp100.com.

```
header-nsrecord, fqdn*, dname*, _new_dname, zone_nameservers*
nsrecord, corp100.com, ns1.corp100.com, ns2.corp100.com, "100.0.0.101/TRUE"
```

This example shows how to overwrite the zone name servers of an existing NS record to 100.0.0.101/TRUE and 100.0.0.102/TRUE.

```
header-nsrecord, fqdn*, dname*, zone_nameservers*
```

```
nsrecord, corp100.com, ns2.corp100.com, "100.0.0.101/TRUE, 100.0.0.102/TRUE"
```

### Merging NS Record Data

This example shows how to merge zone name servers 100.0.0.101/TRUE and 100.0.0.102/TRUE to an NS record.

```
header-nsrecord, fqdn*, dname*, zone_nameservers*
```

```
nsrecord, corp200.com, ns1.corp200.com, "100.0.0.101/TRUE, 100.0.0.102/TRUE"
```

### PTR Record



#### Note

IDN is supported for object type: fqdn and dname. You can use punycode or IDNs while importing these objects.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-PtrRecord	String	Yes			Example: PtrRecord
fqdn	FQDN	No	Name	name zone	This field is required if you do not use the address field. Either the IP address or FQDN is required. Example: 10.0.0.10.in.addr.arpa
_new_fqdn	Reverse FQDN	No			Add this field to overwrite the fqdn field when you use the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default
address	IP address	No	IP Address	ipv4addrss ipv6addrss	This field is required if you do not use the fqdn field. Either the IP address or FQDN is required. Example: 10.0.0.11 If the PTR record belongs to a forward-mapping zone, this field is empty.
_new_address	IP address	No			Add this field to overwrite the address field when you use the overwrite or merge option.
dname	FQDN	Yes	Domain Name	ptrdname	Example: ss.dd.ff
_new_dname	FQDN	No			Add this field to overwrite the dname field when you select the overwrite or merge option.
comment	String	No		Comment comment	
disabled	Boolean	No	Disable	disable	Example: FALSE

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ttl	Unsigned integer	No	TTL	ttl	This is an inherited field. Example: 28800
EA-Site	String	No	Extensible attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for PTR records. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a PTR Record

This example shows how to add a PTR record.

```
header-ptrrecord,dname*,fqdn ptrrecord,ptr.corp100.com,1.0.0.100.in-addr.arpa
```

### Overwriting PTR Record Data

This example shows how to overwrite an existing PTR record with a new FQDN, 2.0.0.100.in-addr.arpa.

```
header-ptrrecord,dname*,fqdn ptrrecord,ptr.corp100.com,2.0.0.100.in-addr.arpa
```

This example shows how to overwrite an existing PTR record with a new IP address, 100.0.0.3.

```
header-ptrrecord,dname*,address ptrrecord,ptr.corp100.com,100.0.0.3
```

### Merging PTR Record Data

This example shows how to change the DNAME of a PTR record from ptr.corp100.com to ptr2.corp100.com, and to add comment = East Asia.

```
header-ptrrecord,dname*,_new_dname,comment
ptrrecord,ptr.corp100.com,ptr2.corp100.com,East Asia
```

## TXT Record



### Note

IDN is supported for object type: fqdn. You can use IDN or punycode while importing this object.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-TxtRecord	String	Yes			Example: TxtRecord
fqdn	FQDN	Yes	Name	name zone	This field combines the record name and the zone name to form the FQDN. Example: t1.corp100.com
_new_fqdn	FQDN	No			Add this field to overwrite the fqdn field when you select the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default
text	String	No	Text	text	You can leave this field empty.
_new_txt	String	No			Add this field to overwrite the txt field when you select the overwrite or merge option.
comment	String	No	Comment	comment	
disabled	Boolean	No	Disable	disable	Example: FALSE
ttl	Unsigned integer	No	TTL	ttl	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 28800
EA-Site	String	No	Extensible attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for TXT record data import. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a TXT Record

This example shows how to add a TXT record, `bind_txt.corp100.com`, with `text = this is a TXT record` and `TTL` set to 3600 seconds.

```
header-txtrecord, fqdn*, text, ttl
```

```
txtrecord, bind_txt.corp100.com, this is a TXT record, 3600
```

## Overwriting TXT Record Data

This example shows how to overwrite the text field of a TXT record.

```
header-txtrecord, fqdn*, text, _new_text
```

```
txtrecord, bind_txt.corp100.com, this is a TXT record, this is a new TXT record
```

## Merging TXT Record Data

This example shows how to add comment = USA and disabled = TRUE to an existing TXT record.

```
header-txtrecord, fqdn*, text, comment, disabled
```

```
txtrecord, bind_txt.corp100.com, this is a TXT record, USA, TRUE
```

## SRV Record



### Note

IDN is supported for object types: fqdn and target. You can use punycode or IDNs while importing these objects.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-SrvRecord	String	Yes			Example: SrvRecord
fqdn	FQDN	Yes	Service	name	This field combines the service name and the zone name to form the FQDN. Example: <i>_http._tcp.corp100.com</i>
_new_fqdn	FQDN	No			Add this field to overwrite the fqdn field when you select the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default priority
priority	Unsigned integer	Yes	Priority	priority	Example: 10
_new_priority	Unsigned integer	No			Add this field to overwrite the priority field when you select the overwrite or merge option.
weight	Unsigned integer	Yes	Weight	weight	Example: 20
_new_weight	Unsigned integer	No			Add this field to overwrite the weight field when you select the overwrite or merge option.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
port	Unsigned integer	Yes	Port	port	Example: 80
_new_port	Unsigned integer	No			Add this field to overwrite the port field when you select the overwrite or merge option.
target	Domain name	Yes	Target	target	Example: foo.test.com
_new_target	Domain name	No			Add this field to overwrite the target field when you select the overwrite or merge option.
comment	String	No	Comment	comment	
disabled	Boolean	No	Disable		disable Example: FALSE
ttl	Unsigned integer	No	TTL	ttl	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 28800
EA-Site	String	No	Extensible attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for SRV records. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a SRV Record

This example shows how to add a new SRV record.

```
header-srvrecord, fqdn*, port*, priority*, target*, weight*
srvrecord, bind_srv.corp100.com, 80, 10, srv.corp100.com, 20
```

## Overwriting SRV Record Data

This example shows how to overwrite the following data of a SRV record: port from 80 to 88 and priority from 10 to 20.

```
header-srvrecord, fqdn*, port*, _new_port, priority*, _new_priority, target*, weight*  
srvrecord, bind_srv.corp100.com, 80, 88, 10, 20, srv.corp100.com, 20
```

This example shows how to overwrite the following data of a SRV record: target from srv.corp100.com to sv2corp100.com and weight from 20 to 30.

```
header-srvrecord, fqdn*, port*, priority*, target*, _new_target, weight*, _new_weight  
srvrecord, bind_srv.corp100.com, 88, 20, srv.corp100.com, sv2.corp100.com, 20, 30
```

## Merging SRV Record Data

This example shows how to merge the following data to a SRV record: comment = USA and disabled = TRUE.

```
header-srvrecord, fqdn*, port*, priority*, target*, weight*, comment, disabled  
srvrecord, bind_srv.corp100.com, 80, 10, srv.corp100.com, 20, USA, TRUE
```

## TLSA Record

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-TlsaRecord	String	Yes			Identifies the first row as a header row for the TLSA record objects. Example: TlsaRecord.
name	String	Yes	Name	name	This field indicates the name of the TLSA resource record. Example: _123._tcp.
certificate_usage	Unsigned integer	Yes	Certificate Usage	certificate_usage	Indicates the data that is used to match the certificate presented in the TLS handshake. Example: 0.
selector	Unsigned integer	Yes	Selector	selector	Specifies which part of the TLS certificate presented by the server is matched with the data during TLS handshake. Example: 1.
matched_type	Unsigned integer	Yes	Matched Type	matched_type	Specifies how the certificate association is presented. Example: 2.
certificate_data	String	Yes	Certificate Data	certificate_data	Either raw data for matching type No hash, or the hash of the raw data for matching types SHA 256 bit and SHA 512 bit. Example: d2abde240d7cd3ee6b4b28c54df034b979 83a1d16e8a410e4561cb106618e971.
_new_certificate_data	String	No	Certificate Data	certificate_data	Add this field to overwrite the certificate_data field when you select the overwrite or merge option.
comment	String	No	Comment	comment	



Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ttl	Boolean	No	TTL	ttl	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> .
fqdn	String	Yes	FQDN	name zone	This field combines the service name and the zone name to form the FQDN. Example: _http_tcp.corp100.com
_new_fqdn	String	No		zone	Add this field to overwrite the fqdn field when you select the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default
disabled	Boolean	No	Disable	disable	Example: FALSE
creator	String	No		creator	Example: STATIC

## Examples

This section contains examples of how to create data files for TLSA records. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a TLSA Record

This example shows how to add a new TLSA record.

```
header-
tlsarecord,certificate_data*,certificate_usage*,fqdn*,matched_type*,selector*,
creator,disabled,name,view,ADMGRP-cloud-api-only,EA-Site
tlsarecord,2D2D2D2D2D424547494E20434552544946494341544520524551554553542D2D2D2D
2D0A4D4
94943354443434163774341514177675A347,0,_443._tcp.doctest,0,0,STATIC,FALSE,_tcp._
443,
default,R0,Boston
```

### Overwriting TLSA Record Data

This example shows how to overwrite the following data of a TLSA record. Consider an example where certificate\_usage\* is 0 as in the following example:

```
header-
tlsarecord,certificate_data*,certificate_usage*,fqdn*,matched_type*,selector*,
creator,disabled,name,view,ADMGRP-cloud-api-only,EA-Site
tlsarecord,2D2D2D2D2D424547494E20434552544946494341544520524551554553542D2D2D2D
2D0A4D4
```

```
94943354443434163774341514177675A347,0,_443._tcp.doctest,0,0,STATIC,FALSE,_tcp._443,
default,R0,Boston
```

Note that `certificate_usage*` is changed from 0 to 1:

```
header-
tlsarecord,certificate_data*,certificate_usage*,fqdn*,matched_type*,selector*,
creator,disabled,name,view,ADMGRP-cloud-api-only,EA-Site
tlsarecord,2D2D2D2D2D424547494E20434552544946494341544520524551554553542D2D2D2D
2D0A4D4
94943354443434163774341514177675A347,1,_443._tcp.doctest,0,0,STATIC,FALSE,_tcp._443,
default,R0,Boston
```

#### Merging TLSA Record Data

This example shows how to disable an existing TLSA record.

```
header-
tlsarecord,certificate_data*,certificate_usage*,fqdn*,matched_type*,selector*,
creator,disabled,name,view,ADMGRP-cloud-api-only,EA-Site
tlsarecord,2D2D2D2D2D424547494E20434552544946494341544520524551554553542D2D2D2D
2D0A4D4
94943354443434163774341514177675A347,0,_443._tcp.doctest,0,0,STATIC,TRUE,_tcp._443,
default,R0,Boston
```

This example shows how to add new FQDN to a TLSA record.

```
header-
tlsarecord,certificate_data*,certificate_usage*,fqdn*,_new_fqdn,matched_type*,
selector*,creator,disabled,name,view,ADMGRP-cloud-api-only,EA-Site
tlsarecord,2D2D2D2D2D424547494E20434552544946494341544520524551554553542D2D2D2D
2D0A4D4
94943354443434163774341514177675A347,0,_443._tcp.doctest,_443._udp.doctest,0,0,STATIC,
FALSE,_tcp._443,default,R0,Boston
```

## CAA Record

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-CAARecord	String	Yes			Identifies the first row as a header row for the CAA record objects. Example: CAARecord.
flag	Int	Yes			Indicates critical or default CAA record. Example: 0
type	String	Yes			Specifies the type of CAA record. Example: issue
CA	String	No			Indicates the name of the certificate authority. Example: certissues.example.com.
ca_details	String	No			Specifies additional details about the CA requests. Example: policy=ev
comment	String	No	Comment	comment	
ttd	Boolean	No	TTL	ttd	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 28800
fqdn	String	Yes	FQDN	name zone	This field combines the service name and the zone name to form the FQDN. Example: _http_tcp.corp100.com
_new_fqdn	String	No		zone	Add this field to overwrite the fqdn field when you select the overwrite or merge option.
disabled	Boolean	No	Disable	disable	Example: FALSE
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default
creator	String	No		creator	Example: STATIC

### Examples

This section contains examples of how to create data files for CAA records. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab. You can add new rows, update existing CAA resource record values or delete an existing CAA resource record.

#### Adding a CAA Record

This example shows how to add a new CAA record.

```
header-
caarecord,ca_value*,_new_ca,ca_flag*,ca_tag*,fqdn*,_new_fqdn,ca_details,_new_ca_
details,comment,creator,disabled,name,ttd,view
```

```
caarecord,abc.com,0,default,issue,xyz.com,0,caa record,new caa record,caa record
information,static,false,0,3600,default
```

### Overwriting CAA Record Data

This example shows how to overwrite the following data of a CAA record. Consider an example where `ca_tag*` is `issue` as in the following example:

```
header-
caarecord,ca_value*,_new_ca,ca_flag*,ca_tag*,fqdn*,_new_fqdn,ca_details,_new_ca_
details,comment,creator,disabled,name,ttl,view
caarecord,abc.com,0,default,issue,xyz.com,0,caa record,new caa record,caa record
information,static,false,0,3600,default
```

Note that `ca_tag*` is changed from `issue` to `issuewild`:

```
header-
caarecord,ca_value*,_new_ca,ca_flag*,ca_tag*,fqdn*,_new_fqdn,ca_details,_new_ca_
details,comment,creator,disabled,name,ttl,view
caarecord,abc.com,0,default,issuewild,xyz.com,0,caa record,new caa record,caa
record information,static,false,0,3600,default
```

### IPv4 Host Address

Use this object type to define parameters for each IP address in an IPv4 host record. The appliance updates an existing host address when you use the "add" mode in an import. For information about host records, refer to the *Infoblox Administrator Guide*.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-HostAddress	String	Yes			Example: HostAddress
parent	FQDN	Yes			Example: h1.corp100.com
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default
network_view	String	No	Network View	network_view	If no network view is specified, the Default view is used. Example: Default
address	IP address	Yes	IP Address	ipv4addrss	Example: 10.0.0.11

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
_new_address	IP address	No			Add this field to overwrite the address field when you select the overwrite or merge option.
mac_address	MAC address	No	MAC Address	mac_address	This is required if the IP address is configured for DHCP. Example: aA:Bb:c2:DD:E1:FF
configure_for_dhcp	Boolean	No	DHCP checkbox	configure_for_dhcp	Example: TRUE
configure_for_dns	integer	No	Enable in DNS	configure for dns	This field is used to specify the parent host record. If not value is specified, TRUE is used. Example: TRUE
deny_bootp	Boolean	No	Deny BOOTP Requests	deny_bootp	Example: FALSE
broadcast_address	String	No	Broadcast Address		
option_logic_filters	List of IPv4 logic filter rules		Filter Type/Action		Examples: .com.infoblox.dns.dhcp_mac_filter\$mac_filter_name, .com.infoblox.dns.nac_filter\$nac_filter_name, .com.infoblox.dns.dhcp_option_filters\$opt_filter_name
boot_file	String	No	Boot File	boot_file	
boot_server	String	No	Boot Server	boot_server	
next_server	String	No	Next Server	next_server	
lease_time	Unsigned integer	No	Lease Time	lease_time	
pxe_lease_time_enabled	Boolean	No	Enable PXE Lease Time		Example: TRUE
pxe_lease_time	Unsigned integer	No	PXE Lease Time	pxe_lease_time	
domain_name	FQDN	No	Domain Name	domain_name	

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
domain_name_servers	IP list	No	Name Server		Example: [5.6.7.8,1.2.3.4]
routers	IP list	No	Router		Example: [2.0.0.2,1.2.3.4]
match_option	String	No		match_client	Data must be in the following format: MAC_ADDRESS/ CLIENT_IDENTIFIER/ RESERVED
ignore_dhcp_param_request_list	Boolean	No	Ignore Optionlist	ignore-dhcp_option_list_request	
OPTION-1	String	No	Custom DHC Options		This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: '255.0.0.0' implies vendor_class='DHCP' (default)
OPTION-XXXX-200	Option information	No	Custom DHC Options		P options This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: 'dfdfdf' name implies vendor_class='XXXX', optioncode/number 200

## Examples

This section contains examples of how to create data files for host addresses. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding an IPv4 Host Address

This example shows how to import the host address 100.0.0.1 in host record h1.corp100.com with a MAC address, enabled DHCP, and a domain name.

```
header-hostaddress,parent*,address*,mac_address,configure_for_dhcp,domain_name
hostaddress,h1.corp100.com,100.0.0.1,aa:aa:aa:aa:aa:aa,TRUE,corp200.com
```

This example shows how to import the host address 100.0.0.2 in host record h2.corp100.com with a MAC address, and two routers with addresses 1.1.1.1 and 2.2.2.2.

```
header-hostaddress,parent*,address*,mac_address,configure_for_dhcp,routers
hostaddress,h2.corp100.com,100.0.0.2,bb:aa:aa:aa:aa:aa,False,1.1.1.1,2.2.2.2
```

## Overwriting IPv4 Host Address Data

This example shows how to overwrite the MAC address and domain name of a host address, and to set configure DHCP to TRUE.

```
header-hostaddress,parent*,address*,mac_address,configure_for_dhcp,domain_name
hostaddress,h1.corp100.com,100.0.0.1,cc:aa:aa:aa:aa:aa,FALSE,corp300.com
```

This example shows how to overwrite the router address of a host address from 1.1.1.1,2.2.2.2 to 1.1.1.1.

```
header-hostaddress,parent*,address*,routers
hostaddress,h2.corp100.com,100.0.0.2,1.1.1.1
```

## Merging IPv4 Host Address Data

This example shows how to merge the following data of a host address: change a new address from 100.0.0.1 to 100.0.0.10, change bootp to DENY, and add lease time.

```
header-hostaddress,parent*,address*,_new_address,deny_bootp,lease_time
hostaddress,h1.corp100.com,100.0.0.1,100.0.0.10,FALSE,3600
```

This example shows how to merge the broadcast address, PXE lease time, and enabled ignore option list to an IPv4 host address.

```
header-
hostaddress,parent*,address*,broadcast_address,pxe_lease_time,ignore_dhcp_param
_request_list hostaddress,h2.corp100.com,100.0.0.2,6.6.6.6,1280,TRUE
```

## IPv6 Host Address

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-IPv6HostAddress	String	Yes			Example: IPv6hostaddress
view	String	No	DNS View	view	If no view is specified, the Default view is used. Example: Default
network_view	String	No	Network View	network_view	If no network view is specified, the Default view is used. Example: Default
parent	FQDN	Yes			Example: h1.corp100.com
address_type	Enumeration	No		address_type	Valid values are ADDRESS, PREFIX, or BOTH. If no value is specified, 'ADDRESS' (default) is used. Example: PREFIX
address	IPv6 address	Yes	IPv6 Address	ipv6addrs	Example: 1001::001
_new_address	IPv6 address	No			Add this field to overwrite the address field when you select the overwrite or merge option.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ipv6_prefix	IPv6 prefix	No		ipv6prefix	This field is required if address_type is 'ADDRESS' or 'BOTH'. Example: 2000:1111::
_new_ipv6_prefix	IPv6 prefix	No			Add this field to overwrite the ipv6_prefix field when you select the overwrite or merge option.
ipv6_prefix_bits	Integer	No		IPv6_prefix_bits	This field is required if address_type is 'ADDRESS' or 'BOTH'. Example: 32
configure_for_dhcp	Boolean	No	DHCP checkbox	configure_for_dhcp	Example: TRUE
configure_for_dns	Boolean	No	Enable in DNS	configure for dns	This field is used to specify the parent host record. If not value is specified, TRUE is used. Example: TRUE
match_option	String	No		match_client	Only 'DUID' is allowed. Example: DUID
duid	String	No	DUID	duid	Example: 0001
domain_name	FQDN	No	Domain Name		
domain_name_servers	IPv6 address list	No	Name Server		Example: '2000::10,3000::10'
valid_lifetime	Unsigned integer	No	Valid Lifetime	valid_lifetime	Example: 43200
preferred_lifetime	Unsigned integer	No	Preferred Lifetime	preferred_lifetime	Example: 604800
OPTION-7	Integer	No	Custom DHCP Options	override_options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: '12' name implies option space = 'DHCPv6', option code/number 7
OPTION-XXXX-200	Option information	No	Custom DHCP Options	override_options	This is an example of a DHCP option. For information, see . Example: 'dfdfdfd' name implies vendor_class='XXXX', option code/number 200

## Examples

This section contains examples of how to create data files for IPv6 host addresses. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.



## Adding an IPv6 Host Address

This example shows how to import an IPv6 host address to a host record in the Default DNS view.

```
header - IPv6hostaddress, parent*, address*, view  
IPv6hostaddress, h1.corp100.com, 1001::001, Default
```

## Overwriting IPv6 Host Address Data

This example shows how to overwrite an IPv6 host address.

```
header - IPv6hostaddress, parent*, address*, _new_address  
IPv6hostaddress, h1.corp100.com, 1001::001, 2000::10
```

## Merging IPv6 Host Address Data

This example shows how to merge a DUID to an IPv6 host address.

```
header - IPv6hostaddress, parent*, address*, duid IPv6hostaddress, h1.corp100.com, 2000::10, 0001
```

## Bulk Host

Field Name	Data Type	Required (Yes/NO)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-BulkHost	String	Yes			Identifies the first row as a header row for the bulk host objects. Example: BulkHost
parentfqdn	FQDN	Yes	Name	zone	Indicates the FQDN of the zone to which bulk host will be added.
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default
prefix	String	No	RFC 2317 Prefix	prefix	Prefix is used for reverse-mapping zones only. The characters must be prepended to the host in the bulk host. If you include a prefix in a forward-mapping zone, the appliance ignores the prefix. No error message is generated.
_new_prefix	String	No			Add this field to overwrite the prefix field when you select the <b>Overwrite</b> or <b>Merge</b> option. Use the hostname of the Grid member in this field. Example: infoblox.localdomain
start_address	IP address	Yes	Start	start_addr	Indicates the starting IP address. Example: 10.0.0.11
_new_start_address	IP address	No			Add this field to overwrite the <b>start_address</b> field when you select the <b>Overwrite</b> or <b>Merge</b> option. Example: 10.0.0.55
end_address	IP address	Yes	End	end_addr	Indicates the last IP address. Example: 10.0.0.22

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
_new_end_address	IP address	No			Add this field to overwrite the <b>end_address</b> field when you select the <b>Overwrite</b> or <b>Merge</b> option. Example: 10.0.0.66
reverse	Boolean	No			Example: True
comment	String	No	Comment	comment	Example: This is a Bulk Host.
disabled	Boolean	No	Disable	disable	Enable or disable the bulk host. Example: FALSE
ttl	Unsigned integer	No	TTL	ttl	This is an inherited field. Example: 7200
EA-Site	String	No	Extensible attribute	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. Example: California.
EA-Users	List	No	Extensible attribute	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. Example: ['Annie', 'John'].
ADMGRP-XXXX	String	No	Permissions Admin Group/ Role	permission	ADMGRP-JimSmith is an example of an admin permission of a specific admin group. Example: RW

## NXDOMAIN and Blacklist Rulesets

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-Ruleset	String	Yes			Example: Ruleset
name	String	Yes	Name	name	Example: ruleset1
_new_name	String	No			Example: ruleset1-new
type	String	Yes		type	You can use this field for NXDOMAIN rules or blacklist rules. Valid value is NXDOMAIN or BLACKLIST. Example: NXDOMAIN
disabled	Boolean	No	Disable	disabled	Example: FALSE
comment	String	No	Comment	comment	
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for NXDOMAIN and blacklist rulesets. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding an NXDOMAIN Ruleset

This example shows how to import an NXDOMAIN ruleset.

```
header-Ruleset,name*,type*,comment,disabled ruleset,NXD,NXDOMAIN,This is an  
NXDOMAIN ruleset,FALSE
```

### Adding a Blacklist Ruleset

This example shows how to import a blacklist ruleset.

```
header-Ruleset,name*,type*,disabled,comment  
ruleset,blacklistrule1,BLACKLIST,FALSE,This is a blackset ruleset
```

### Overwriting Blacklist Ruleset Data

This example shows how to overwrite the name of a blacklist ruleset.

```
header-Ruleset,name*,_new_name,type*  
ruleSet,blacklistrule1,blacklistrule2,BLACKLIST
```

### Merging NXDOMAIN Ruleset Data

This example shows how to merge an admin permission to a NXDOMAIN ruleset.

```
header-Ruleset,name*,type*,ADMGRP-JoeSmith ruleSet,NXD,NXDOMAIN,RW
```

## NXDOMAIN Rule

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-nxdomainrule	String	Yes			Example: NxdomainRule
parent	String	Yes	Name	nxdomain_rules	This field indicates the ruleset to which the NXDOMAIN rule belongs. Example: nxdomain-ruleset1
pattern	String	Yes	Pattern	pattern	Use this to match domain names. You cannot use the characters * and ,, in the domain name. Example: foo
_new_pattern	String	No			Example: foo-new

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
action	String	Yes	Action	action	The valid value is PASS, MODIFY or REDIRECT. Example: PASS

## Examples

This section contains examples of how to create data files for NXDOMAIN rules. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding an NXDOMAIN Rule

This example shows how to import an NXDOMAIN rule.

```
header-nxdomainrule,action*,parent*,pattern* nxdomainrule,REDIRECT,NXD,*foo.com
```

### Overwriting NXDOMAIN Rule Data

This example shows how to overwrite the action of an NXDOMAIN rule.

```
header-nxdomainrule,action*,parent*,pattern* nxdomainRule,REDIRCT,NXD,*bar.com
```

## Blacklist Rule

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-blacklistrule	String	Yes			Example: BlacklistRule
parent	String	Yes	Name	blacklist_rulesets	Example: blacklist-ruleset1
domain_name	String	Yes	Domain Name		Use this to match domain names. You cannot use the characters * and , in the domain name. The domain name cannot exceed 255 characters. Example: www.foo.com
_new_domain_name	String	No			Example: www.bar.com
action	String	Yes	Action	blacklist_action	The valid value is PASS or REDIRECT. Example: PASS

## Examples

This section contains examples of how to create data files for blacklist rules. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a Blacklist Rule

This example shows how to import a blacklist rule.

```
header-blacklistrule,parent*,domain_name*,action*
```

```
BlacklistRule,BlackList,foobar.com,REDIRECT
```

### Overwriting Blacklist Rule Data

This example shows how to overwrite the action of a blacklist rule.

```
header-blacklistrule,parent*,domain_name*,action*
```

```
BlacklistRule,BlackList,foobar.com,PASS
```

## Whitelist Rule

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
header-analyticsdomainwhitelist	String	Yes			Example: analyticsdomainwhitelist
fqdn	String	Yes	Domain Name		Use this to match domain names. You cannot use the characters * and , in the domain name. The domain name cannot exceed 255 characters. Example: www.test.com
comment	String	No	Comment		Enter additional information about this domain.
disabled	String	Yes	Disable		The valid value is TRUE or FALSE. Example: FALSE
type	CUSTOM	Yes			Example: CUSTOM

### Example

This section contains an example of how to create data file for whitelist rules. The example uses comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a Whitelist Rule

This example shows how to import a whitelist rule.

```
header-analyticsdomainwhitelist, fqdn*, disabled, type
```

```
BlacklistRule,BlackList,foobar.com,FALSE,CUSTOM
```

## DNS64 Synthesis Group

You can import the name of a named ACL in the fields that support named ACLs, such as clients, mapped, and exclude.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-Dns64SynthesisGroup	String	Yes			Example: Dns64SynthesisGroup
name	String	Yes	Name	name	Example: group1
prefix	IPv6 network	Yes	Prefix	prefix	Example: 64:FF9B::/96
comment	String	No	Comment	comment	
clients	Access control list	No	Name	clients	Valid values are IPv4 and IPv6 addresses and networks only. The default value is 'Any'. Example: 2000::/64/AllowNote that you can import the name of a named ACL in this field.
mapped	Access control list	No	Mapped IPv4 Address → Name	mapped	Valid values are IPv4 addresses and networks only. The default value is 'Any'. Example: 10.0.0.0/8/AllowNote that you can import the name of a named ACL in this field.
exclude	Access control list	No	Excluded IPv6 Address → Name	exclude	Valid values are IPv6 addresses and networks only. The default is 'None'. Example: 2000::/AllowNote that you can import the name of a named ACL in this field.
disabled	Boolean	No	Disabled	disable	Example: FALSE

## Examples

This section contains examples of how to create data files for DNS64 synthesis groups. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a DNS64 Synthesis Group

This example shows how to import a DNS64 Synthesis Group.

```
header-Dns64SynthesisGroup,name*,prefix*,mapped
Dns64SynthesisGroup,DNS64Group1,64:FF9B::/96,10.0.0.0/8/Allow
```

### Overwriting DNS64 Synthesis Group Data

This example shows how to overwrite the mapped IPv4 address of a DNS64 synthesis group.

```
header-Dns64SynthesisGroup,name*,prefix*,mapped
Dns64SynthesisGroup,DNS64Group1,64:FF9B::/96,10.1.1.0/24/Allow
```

### Adding Named ACL Data

This example shows how to add the name of a named ACL to the mapped address of a DNS64 synthesis group.

header=Dns64SynthesisGroup,name\*,prefix\*,mapped

Dns64SynthesisGroup,DNS64Group1,My\_Named\_ACL

## Response Policy Zone

For a Response Policy Zone, CSV import supports all the values that are listed in [Authoritative Zone](#) along with the following values:

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
priority	Unsigned integer	No			Example: 1001
severity	String	No	Severity		Valid values are Critical, Major, Warning, and Informational
substitute_name	FQDN	No			Example: JKL.NET
rpz_policy	String	No	Policy Override		Valid values are GIVEN, NXDOMAIN, NODATA, PASSTHRU, SUBSTITUTE, and DISABLED
rpz_drop_ip_rule_enabled	Boolean	No	Ignore RPZ-IP triggers with too small prefix lengths		When this is set to <b>True</b> , DNS server ignores RPZ-IP rules with prefix lengths that are less than the specified prefix length limit. Example: TRUE
rpz_drop_ip_rule_min_prefix_length_ipv4	Unsigned Integer	No	Minimum IPv4 Prefix Length		Indicates the minimum IPv4 prefix length for RPZ-IP triggers. The default value is 29.
rpz_drop_ip_rule_min_prefix_length_ipv6	Unsigned Integer	No	Maximum IPv4 Prefix Length		Indicates the minimum IPv6 prefix length for RPZ-IP triggers. The default value is 112.

## Response Policy A Record

For a Response Policy A Record, CSV import supports all the values that are listed in [A Record](#) along with the following values:

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
create_ptr	Boolean	No	Create associated PTR record		Example: TRUE
parent_zone	FQDN	No			Example: ABC.NET

## Response Policy AAAA Record

For a Response Policy AAAA Record, CSV import supports all the values that are listed in [AAAA Record](#) along with the following value:

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
parent_zone	FQDN	No			Example: ABC.NET

## Response Policy IP A Record

For a Response Policy IP A Record, CSV import supports all the values that are listed in [A Record](#) along with the following values:

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
create_ptr	Boolean	No			Example: TRUE
parent_zone	FQDN	No			Example: ABC.NET

## Response Policy IP AAAA Record

For a Response Policy IP AAAA Record, CSV import supports all the values that are listed in [AAAA Record](#) along with the following value:

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
parent_zone	FQDN	No			Example: ABC.NET

## Response Policy MX Record

For a Response Policy MX Record, CSV import supports all the values that are listed in [MX Record](#) along with the following value:

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
parent_zone	FQDN	No			Example: ABC.NET

## Response Policy NAPTR Record

For a Response Policy NAPTR Record, CSV import supports all the values that are listed in [NAPTR Record](#) along with the following value:



Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
parent_zone	FQDN	No			Example: ABC.NET

## Response Policy PTR Record

For a Response Policy PTR Record, CSV import supports all the values that are listed in [PTR Record](#) along with the following value:

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
parent_zone	FQDN	No			Example: ABC.NET

## Response Policy SRV Record

For a Response Policy SRV Record, CSV import supports all the values that are listed in [SRV Record](#) along with the following value:

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
parent_zone	FQDN	No			Example: ABC.NET

## Response Policy TXT Record

For Response Policy TXT Record, CSV import supports all the values that are listed in [TXT Record](#) along with the following value:

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
parent_zone	FQDN	No			Example: ABC.NET

## Response Policy CNAME Record

For Response Policy CNAME Record, CSV import supports all the values that are listed in [CNAME Record](#) along with the following value:

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
parent_zone	FQDN	No			Example: ABC.NET

## Response Policy IP Address

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-ResponsePolicyIPAddress	String	Yes			Example: ResponsePolicyIPAddress
fqdn	FQDN	Yes	Name	name	Example: 10.1.2.3.ABC.NET
_new_fqdn	FQDN	No			Add this field to overwrite the fqdn field when you select the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used.
canonical_name	IP address	No	Canonical Name	canonical	Example: 10.1.2.3
parent_zone	FQDN	No			Example: ABC.NET
comment	String	No	Comment	comment	
disabled	Boolean	No	Disable	disable	Example: FALSE
ttl	Unsigned integer	No	TTL	ttl	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 7200

### Example

This example shows how to create data files for Response Policy IP Address. The example uses comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

```
header-ResponsePolicyIPAddress,fqdn*,canonical_name,disabled,parent_zone,view
ResponsePolicyIPAddress,10.1.2.3.ABC.NET,10.1.2.3,FALSE,ABC.NET,DEFAULT
```

## Response Policy Client IP Address

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-ResponsePolicyClientIPAddress	String	Yes			Example: ResponsePolicyClientIPAddress
fqdn	FQDN	Yes	Name	name	Example: 10.1.2.1.ABC.NET
_new_fqdn	FQDN	No			Add this field to overwrite the fqdn field when you select the overwrite or merge option.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
view	String	No	DNS View	views	If no view is specified, the Default view is used.
canonical_name	IP address	No	Canonical Name	canonical	Example: 10.1.2.1
parent_zone	FQDN	No			Example: ABC.NET
comment	String	No	Comment	comment	
disabled	Boolean	No	Disable	disable	Example: FALSE
ttl	Unsigned integer	No	TTL	ttl	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 7200

### Example

This example shows how to create data files for Response Policy Client IP Address. The example uses comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

header-

```
ResponsePolicyClientIPAddress, fqdn*, canonical_name, disabled, parent_zone, view
ResponsePolicyClientIPAddress, 10.1.2.1.ABC.NET, 10.1.2.1, FALSE, ABC.NET, DEFAULT
```

### Response Policy IP Address CNAME

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-ResponsePolicyIPAddressCname	String	Yes			Example: ResponsePolicyIPAddressCname
fqdn	FQDN	Yes	Name	name	Example: 10.1.2.3.ABC.NET
_new_fqdn	FQDN	No			Add this field to overwrite the fqdn field when you select the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used.
canonical_name	IP address	No	Canonical Name	canonical	Example: 10.1.2.3
parent_zone	FQDN	No			Example: ABC.NET
comment	String	No	Comment	comment	

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
disabled	Boolean	No	Disable	disable	Example: FALSE
ttl	Unsigned integer	No	TTL	ttl	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 7200

### Example

This example shows how to create data files for Response Policy IP Address CNAME. The example uses comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

header-

```
ResponsePolicyIPAddressCname, fqdn*, canonical_name, disabled, parent_zone, view
ResponsePolicyIPAddressCname, 10.1.2.3.ABC.NET, 10.1.2.3, FALSE, ABC.NET, DEFAULT
```

### Response Policy Client IP Address CNAME

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-ResponsePolicyClientIPAddress	String	Yes			Example: ResponsePolicyClientIPAddress
fqdn	FQDN	Yes	Name	name	Example: 10.1.2.1.ABC.NET
_new_fqdn	FQDN	No			Add this field to overwrite the fqdn field when you select the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used.
canonical_name	IP address	No	Canonical Name	canonical	Example: 10.1.2.1
parent_zone	FQDN	No			Example: ABC.NET
comment	String	No	Comment	comment	
disabled	Boolean	No	Disable	disable	Example: FALSE
ttl	Unsigned integer	No	TTL	ttl	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 7200

### Example

This example shows how to create data files for Response Policy Client IP Address CNAME. The example uses comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

header-

ResponsePolicyClientIPAddressCname, fqdn\*, canonical\_name, disabled, parent\_zone, view

ResponsePolicyClientIPAddressCname, 10.1.2.1.ABC.NET, 10.1.2.1, FALSE, ABC.NET, DEFAULT

## Network View

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header- NetworkView	String	Yes			Identifies the first row as a header row for the network view objects. Example: NetworkView
name	String	Yes			Indicates the name of the network view. Example: net_view1
_new_name	String	No			Add this field to overwrite the <b>name</b> field when you select the <b>Override</b> or <b>Merge</b> option.
comment	String	No	Comment	comment	Example: This is a network view.
EA-Site	String	No	Extensible attribute	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. Example: California.
EA-Users	String	No	Extensible attribute	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. Example: John.
ADMGRP-XXXX	String	No	Permissions Admin Group/Role	permission	ADMGRP-JimSmith is an example of an admin permission of a specific admin group. Example: RW

## DNS View

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guideline
Header-View	String	Yes			Identifies the first row as a header row for the DNS view objects. Example: View
name	String	Yes			Example: dns_view1

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guideline
_new_name	String	No			Add this field to overwrite the <b>name</b> field when you select the <b>Overwrite</b> or <b>Merge</b> option.
comment	String	No	Comment	comment	Example: This is a DNS view.
network_view	String	No	Network View	network_view	If no network view is specified, the default view is used. Example: Default
disable	Boolean	No	Disable	disable	Enable or disable view. Example: False
recursion	Boolean	No		allow_recursive_query	Flag to respond to recursive queries. Example: False
root_name_server_type	String	No			This is a single inheritance from GridDns. Example: Custom
match_clients	ACL	No			List of <b>address_tsig_ac</b> items. It can be an IP address, a network entry, <b>Any</b> or a <b>TSIG-</b> /permission. If the first value is not <b>Any</b> or <b>TSIG-</b> , it is assumed to be an IP address or a network entry. Example: 10.0.0.10/Allow, 11.0.0.0/16/Deny, TSIG-foo/xyz/Allow. It can also be a named ACL. Example: NACL1.
match_destinations	ACL	No			Indicates the match_destination list. It can be an IP address, a network, <b>Any</b> or a <b>TSIG-</b> /permission. If the first value is not <b>Any</b> or <b>TSIG-</b> , it is assumed to be an IP address or a network entry. Example: 10.0.0.10/Allow, 11.0.0.0/16/Deny, TSIG-foo/xyz/Allow. It can also be a named ACL. Example: NACL1.
custom_root_name_servers	Root nameserver list	No			List of custom root name servers. Example: mm1.test.com/1.1.1.1/
lame_ttl	Unsigned integer	No			Indicates the lame TTL value in seconds. Example: 600
nxdomain_redirect	Boolean	No			Enable intercept and redirect nxdomain responses Example: False

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guideline
nxdomain_redirect_addresses	IP address list	No			Example: 1.1.1.1,2.2.2.2 Note that the field <b>nxdomain_redirect_addresses</b> is dependent on the <b>nxdomain_redirect</b> field. Infoblox recommends that you specify values for <b>nxdomain_redirect_addresses</b> and <b>nxdomain_redirect</b> fields and do not leave these fields blank while performing a CSV import operation.
nxdomain_redirect_ttl	Unsigned integer	No			Indicates the NXDOMAIN redirect ttl in seconds. Example: 60
nxdomain_log_query	Boolean	No			When you set this to <b>True</b> , NXDOMAIN redirections will be logged. Example: False
nxdomain_rulesets	List of domain names	No			List of Ruleset objects used for NXDOMAIN redirection. Example: nxd1.com, nxd2.com
enable_blacklist	Boolean	No		enable_blacklist	Enable or disable blacklisting at the Grid level. Example: False
blacklist_redirect_addresses	IP address list	No		blacklist_redirect_addresses	Set or retrieve the list of IPv4 addresses to which the blacklisted queries are redirected. Example: 1.1.1.1,2.2.2.2
blacklist_action	String	No			Specify the action to be performed when a domain name matches the pattern defined in an assigned rule. Example: Redirect
blacklist_redirect_ttl	Unsigned integer	No			Set or retrieve the TTL value of synthetic DNS responses resulted by blacklisted queries. Example: 60
blacklist_log_query	Boolean	No		blacklist_log_query	Specify if blacklisted queries must be logged. Example: False
blacklist_rulesets	List of domain names	No		blacklist_rulesets	Specify or retrieve ruleset objects that are blacklisted at the Grid level. Example: list1.com, list2.com, ...
enable_dns64	Boolean	No			Enable or disable DNS64 synthesis. Example: False
dns64_groups	DNS64 group list	No			List of SynthesisGroup objects. Example: dns64_grp1, dns64_grp2
forwarders_only	Boolean	No	Use Forwarders Only	forward_only	Enable use of forwarders only. Example: False

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guideline
forwarders	IP address list	No		forwarders	List of forwarders for zone transfers. Example: "10.10.0.1,20.20.0.1,.."
filter_aaaa	String	No			Indicates the type of AAAA filtering for this Grid DNS object. The default value is <b>No</b> . Example: Yes
filter_aaaa_list	ACL	No			Indicates the list of IPv4 addresses and networks from which queries are received. Note that the AAAA filtering is applied to these addresses. Example: "12.0.0.12/ Deny,13.0.0.0/8/Allow.. ." or "NACL1"
max_cache_ttl	Unsigned integer	No			Indicates the maximum time (in seconds) for which the server will cache positive answers. The default value is 604800.
max_nocache_ttl	Unsigned integer	No			Indicates the maximum time (in seconds) for which the server will cache negative (NXDOMAIN) responses. The default value is 10800. The maximum allowed value is 604800.
EA-Site	String	No	Extensible attribute	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. Example: California.
EA-Users	String	No	Extensible attribute	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. Example: John.
ADMGRP-XXXX	String	No	Permissions Admin Group/Role	permission	ADMGRP-JimSmith is an example of an admin permission of a specific admin group. Example: RW
rpz_drop_ip_rule_enabled	Boolean	No	Ignore RPZ-IP triggers with too small prefix lengths		When this is set to True, DNS server ignores RPZ-IP rules with prefix lengths that are less than the specified prefix length limit. Example: TRUE
rpz_drop_ip_rule_min_prefix_length_ip v4	Unsigned Integer	No	Minimum IPv4 Prefix Length		Indicates the minimum IPv4 prefix length for RPZ-IP triggers. The default value is 29.
rpz_drop_ip_rule_min_prefix_length_ip v6	Unsigned Integer	No	Maximum IPv4 Prefix Length		Indicates the minimum IPv6 prefix length for RPZ-IP triggers. The default value is 112



## IPv4 Network Container

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guideline
Header- NetworkContainer	String	Yes			Identifies the first row as a header row for network container objects. Example: NetworkContainer
address	IP address	Yes	IP Address	ipv4addr	Indicates the IP address of the network container. Example: 192.138.1.1
netmask	Prefix	Yes	Netmask	network	Indicates the subnet mask of a network container in CIDR format. Example: 24
comment	String	No	Comment	comment	Example: This is an IPv4 network container.
lease_time	integer	No	Lease Time		Example: 1100
routers	IP address list	No	Routers		Example: "10.0.0.1,10.0.0.100,"
domain_name	FQDN	No	Domain Name		
domain_name_servers	IP address list	No	DNS Servers		Example: "10.2.3.4,11.2.3.4"
broadcast_address	IP address	No	Broadcast Address		Example: 10.0.0.1
OPTION-1	String	No	Custom DHCP Options	options	This is an example of a DHCP option. For Options information, see <a href="#">Data Specific Guidelines</a> . Example: '255.0.0.0' name implies vendor_class='DHCP' (default)
OPTION-XXXX-200	Option information	No	Custom DHCP Options	options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: 'dfdfdfd' name implies vendor_class='XXXX', option code/number 200
enable_ddns	Boolean	No	Enable DDNS Updates	enable_ddns	Example: FALSE
ddns_domainname	String	No	DDNS Domain Name	ddns_domainname	Example: ddns.corp100.com

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guideline
ddns_ttl	Unsigned integer	No	DDNS Update TTL	ddns_ttl	Indicates the DDNS TTL value in seconds. Example: 1200
ddns_generate_hostname	Boolean	No	Generate Hostname	ddns_generate_hostname	When you set this to <b>True</b> , DHCP server will generate a hostname for DNS updates if not sent by client. Example: TRUE
update_static_leases	Boolean	No	Update Fixed Address	ddns_update_fixed_addresses	Example: FALSE
enable_option81	Boolean	No	Option 81 Support	ddns_use_option81	Example: TRUE
update_dns_on_lease_renewal	Boolean	No	Lease Renewal Update	override_update_dns_on_lease_renewal	Example: TRUE
enable_dhcp_thresholds	Boolean	No	Enable DHCP Thresholds	enable_dhcp_thresholds	When you set this field to TRUE, you must enter values in the range_high_water_mark and range_low_water_mark fields. You cannot leave those fields empty. Otherwise, the appliance generates an error.
enable_email_warnings	Boolean	No		enable_email_warnings	Enable to send DHCP threshold warnings via email. Example: False
enable_snmp_warnings	Boolean	No			Enable to send DHCP threshold warnings via SNMP. Example: False
threshold_email_addresses	email address list	No	Email Addresses		Example: "admin1@infoblox.com';admin2@somewhere.com"
pxe_lease_time	Unsigned integer	No	PXE Lease Time	pxe_lease_time	Example: 1100
deny_bootp	Boolean	No	Deny BOOTP Requests	deny_bootp	Example: FALSE
boot_file	String	No	Boot File	bootfile	Example: bootfile1
boot_server	String	No	Boot Server	bootserver	Example: abc.corp100.com
next_server	String	No	Next Server	nextserver	Example: blue.domain.com

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guideline
option_logic_filters	List of IPv4 logic filter rules		Filter Type/ Action		Examples: .com.infoblox.dns.dhcp_mac_filter\$mac_filter_name,.com.infoblox.dns.nac_filter\$nac_filter_name,.com.infoblox.dns.dhcp_option_filters\$opt_filter_name
lease_scavenge_time	Unsigned integer	No			Indicates the Grid level <b>lease_scavenge_time</b> value. If the value is -1, which means this lease scavenge will be disabled. The minimum value would be 7 * 24 * 60 * 60 ( 7 days).
is_authoritative	Boolean	No	Authoritative	authority	Example: FALSE
recycle_leases	Boolean	No	Lease Deletion	recycle_leases	This field is set to TRUE by default. Ensure that you use the overwrite option if you want to change the value to FALSE. Merging data from an import preserves the default value.
ignore_client_requested_options	Boolean	No	Ignore Optionlist	ignore_dhcp_option_list_request	Example: TRUE
network_view	String	No	Network View	network_view	If no network view is specified, the default view is used. Example: Default
rir_organization	String	No	Organization Name	rir_organization	Identifies the Regional Internet Registry (RIR) organization object. Use this only when the network is associated with an RIR organization. Example: corp100
rir_registration_status	String	No	Registration Status	rir_registration_status	Identifies the registration status of Regional Internet Registry (RIR). Use this only for an RIR network. When you enable the <b>Enable Updates Of RIR Registrations</b> checkbox at the Grid level and import a CSV file to add either an <b>IPv4 network container</b> or an <b>IPv6 network container</b> with the <b>rir_registration_status</b> set to Registered without values for any other RIR fields, the appliance completes the import operation and adds the IPv4 network container or the IPv6 network container to the Grid. The status of this IPv4 network container or the IPv6 network container is set as <b>Non-registered</b> network. Example: Non-registered
last_rir_registration_update_sent	String	No			Identifies the last registration update timestamp of Regional Internet Registry (RIR). This is a read-only attribute.
last_rir_registration_update_status	String	No			Identifies the last registration update status of Regional Internet Registry (RIR). This is a read-only attribute.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guideline
enable_discovery	Boolean	Yes	Enable Discovery	network	If this field is set to <b>True</b> , the <b>discovery_member</b> must also be defined.
discovery_member	String	Yes	Discovery Member	network	Indicates the discovery member name. Required if discovery is enabled for the network.
discovery_exclusion_range	Prefix	No			List of IP ranges to be excluded from the discovery process.
remove-subnets	Boolean	No			Specify <b>False</b> to keep the subnets or <b>True</b> to remove them. The default value is undefined, which is to remove all subnets. Use this only when you want to delete a network container. When you perform CSV export of a network container, the CSV file does not contain the <b>remove-subnets</b> column. You must add the <b>remove-subnets</b> column to the CSV export file if you want to perform a delete operation. When you delete a network container from the CSV file, you must set <b>remove-subnets</b> to <b>True</b> to delete both the parent and sub networks. When you set this column value to <b>False</b> , the CSV delete operation deletes only the parent network when you delete a network container.
EA-Site	String	No	Extensible attribute	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. Example: California.
EA-Users	String	No	Extensible attribute	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. Example: John.
ADMGRP-XXXX	String	No	Permissions Admin Group/ Role	permission	ADMGRP-JimSmith is an example of an admin permission of a specific admin group. Example: RW

## IPv4 Network

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-Network	String	Yes			Example: Network
rir_organization	String	No	Organization Name	rir_organization	Use this only when the network is associated with an RIR organization. Example: corp100

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
rir_registration_status	String	No	Registration Status	rir_registration_status	Use this only when this is an RIR network. Example: Registered
address	IP address	Yes	Address	network	Example: 10.0.0.11
netmask	Netmask	Yes	Netmask		Example: 255.255.0.0
network_view	String	No	Network View	network_view	If no view is specified, the Default view is used. Example: Default
enable_discovery	Boolean	Yes	Enable Discovery	enable_discovery	If this field is set to TRUE, the discovery_member must also be defined.
discovery_member	String	Yes	Discovery Member	discovery_member	Required if discovery is enabled for the network.
discovery_exclusion_range	IP Prefix	No	Network Editor → Discovery Exclusions	discovery_exclusion_range	List of IP ranges to be excluded from the discovery process.
comment	String	No	Comment	comment	
auto_create_reversezone	Boolean	No	Automatically create reverse mapping zone	auto_create_reversezone	Example: TRUE
is_authoritative	Boolean	No	Authoritative	authority	Example: FALSE
option_logic_filters	List of IPv4 logic filter rules	No	Filter Type/Action		Examples: .com.infoblox.dns.dhcp_mac_filter\$mac_filter_name, .com.infoblox.dns.nac_filter\$nac_filter_name, .com.infoblox.dns.dhcp_option_filters\$opt_filter_name
boot_file	String	No	Boot File	bootfile	Example: bootfile1
boot_server	String	No	Boot Server	bootserver	Example: <a href="#">abc.corp100.com</a>
ddns_domainname	String	No	DDNS Domain Name	ddns_domainname	Example: <a href="#">ddns.corp100.com</a>
generate_hostname	Boolean	No	Generate Hostname	generate_hostname	When you set this to <b>True</b> , DHCP server will generate a hostname for DNS updates if not sent by client. Example: TRUE
always_update_dns	Boolean	No	DNS Zones Associations	ddns_server_always_updates	Example: FALSE

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
update_static_leases	Boolean	No	Fixed Address Updates	ddns_update_fixed_address	Example: FALSE
update_dns_on_lease_renewal	Boolean	No	Update DNS on DHCP Lease Renewal	override_update_dns_on_lease_renewal	Example: TRUE
ddns_ttl	Unsigned integer	No	DDNS Update TTL	ddns_ttl	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 1200
enable_option81	Boolean	No	Option 81 Support	ddns_use_option81	Example: TRUE
deny_bootp	Boolean	No	Deny BOOTP Requests	deny_bootp	Example: FALSE
broadcast_address	String	No	Broadcast Address		
disabled	Boolean	No	Disabled	disable	Example: FALSE
enable_ddns	Boolean	No	Enable DDNS Updates	enable_ddns	Example: FALSE
enable_thresholds	Boolean	No	Enable DHCP Thresholds	enable_dhcp_thresholds	When you set this field to TRUE, you must enter values in the range_high_water_mark and range_low_water_mark fields. You cannot leave those fields empty. Otherwise, the appliance generates an error.
enable_threshold_email_warnings	Boolean	No	Enable Email Warnings	enable_email_warnings	When you use the merge function, the appliance preserves the existing value in this field. When you use the overwrite function, you must include a value (TRUE or FALSE). Otherwise, the appliance generates an error. Example: TRUE
enable_threshold_snmp_warnings	Boolean	No	Enable SNMP Warnings	enable_snmp_warnings	When you use the merge function, the appliance preserves the existing value in this field. When you use the overwrite function, you must include a value (TRUE or FALSE). Otherwise, the appliance generates an error. Example: TRUE
range_high_water_mark	Integer	No	High Water Mark	high_water_mark	When you set enable_thresholds to TRUE, you must enter values in this field and in the range_low_water_mark field. You cannot leave these fields empty. Otherwise, the appliance generates an error. Example: 80

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ignore_client_requested_options	Boolean	No	Ignore Optionlist	ignore_dhcp_option_list_request	Example: TRUE
range_low_water_mark	Integer	No	Low Water Mark	low_water_mark	When you set enable_thresholds to TRUE, you must enter values in this field and in the range_high_water_mark field. You cannot leave these fields empty. Otherwise, the appliance generates an error. Example: 10
next_server	String	No	Next Server	nextserver	Example: <a href="#">blue.domain.com</a>
lease_time	Unsigned integer	No	Lease Time		Example: 1100
enable_pxe_lease_time	Boolean	No	Enable PXE lease time		Example: FALSE
pxe_lease_time	Unsigned integer	No	PXE Lease Time	pxe_lease_time	Example: 1100
recycle_leases	Boolean	No	Lease Deletion	recycle_leases	This field is set to TRUE by default. Ensure that you use the overwrite option if you want to change the value to FALSE. Merging data from an import preserves the default value.
threshold_email_addresses	email address list	No	Email Addresses		Example: "admin1@infoblox.com', 'admin2@somewhere.com"
dhcp_members	Grid member list	No	Members	members	Example: <a href="#">"host1.infoblox.com, host2.infoblox.com"</a>
routers	IP address list	No	Routers		Example: "10.0.0.1,10.0.0.100,"
domain_name	FQDN	No	Domain Name		
domain_name_servers	IP address list	No	DNS Servers		Example: "10.2.3.4,11.2.3.4"
zone_associations	Zone list	No			Example: <a href="#">test.com/TRUE</a>
VLANs	String	No	Assigned VLAN ID Assigned VLAN Name		VLAN View is a container object which can contain VLAN Range/VLAN objects. Example: default/1/4094/1

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
OPTION-1	String	No	Custom DHCP Options	options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: '255.0.0.0' name implies vendor_class='DHCP' (default)
OPTION-XXXX-200	Option information	No	Custom DHCP Options	options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: 'dfdfdf' name implies vendor_class='XXXX', option code/number 200
EA-Site	String	No	Extensible Attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-XXX	String	No	Extensible Attribute XXX value	inheritable	EA-XXX is an example of an inheritable extensible attribute where XXX represents the attribute name. This column displays the value of an inheritable extensible attribute. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EAIherited-XXX	String	No	Inheritance State of an Extensible Attribute XXX	inheritable	EAIherited-XXX is an example of an inheritable extensible attribute where XXX represents the attribute name. This column displays the inheritance state of an inheritable extensible attribute. This column is displayed only if the extensible attribute is inheritable. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Extensible Attribute User	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for networks. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.



## Adding an IPv4 Network

This example shows how to import network 10.1.0.0/16 in the network view External with auto create reverse zone enabled and a comment.

```
header-network,address*,netmask*,network_view,auto_create_reversezone,comment
network,10.1.0.0,255.255.0.0,External,TRUE,This is comment field
```

For a network that has discovery enabled:

```
header-network,address*,netmask*,network_view,auto_create_reversezone,
enable_discovery,discovery_member
network,10.1.0.0,255.255.0.0,External,TRUE,TRUE,gridmember1.localdomain
```

For a network that has discovery enabled and including an IPv4 range that is excluded for discovery:

```
header-network,address*,netmask*,network_view,auto_create_reversezone,
enable_discovery,discovery_member,discovery_exclusion_range
network,10.1.0.0,255.255.0.0,External,TRUE,gridmember1.localdomain,10.1.8.0-10.1
.8.255
```

You cannot import a network container, but you can add an IPv4 network container through GUI.

## Overwriting IPv4 Network Data

This example shows how to overwrite the following data in an existing network: enable the network to be "authoritative" and to add boot files bppt\_file\_001 and boot server 1.2.3.4.

```
header-network,address*,netmask*,is_authoritative,boot_file,boot_server
network,100.0.0.0,255.255.255.0,True,boot_file_001,1.2.3.4
```

## Merging IPv4 Network Data

This example shows how to merge the following data to an existing network: DDNS domain name, enable generate hostname, and disable always update DNS.

```
header-
network,address*,netmask*,ddns_domainname,generate_hostname,always_update_dns
network,100.0.0.0,255.255.255.0,ddns.corp100.com,TRUE,FALSE
```

## Adding IPv4 Networks with Zone Associations Enabled

This example shows how to add networks to a member in the default network view with zone association added as the default. The first row adds a network "20.0.1.0/24" that maps to member "[ib-10-34-43-2.infoblox.com](#)" in the "default" network view with zone association "[dnszone1.com](#)" added as the default zone association. The second row adds a network "20.0.2.0/24" that maps to member "[ib-10-34-43-2.infoblox.com](#)" in the "default" network view with zone associations "[dnszone1.com](#)" and "[dnszone2.com](#)" added, where zone "[dnszone2.com](#)" is configured as the default zone association.

```
network,address*,netmask*,dhcp_members,network_view,zone_associations
```

```
network,20.0.1.0,255.255.255.0,ib-10-34-43-2.infoblox.com,default,dnszone1.com/True/default
```

```
network,20.0.2.0,255.255.255.0,ib-10-34-43-2.infoblox.com,default,dnszone1.com/False/default, dnszone2.com/True/default
```

### Overwriting IPv4 Network Data with Zone Associations Enabled

This example shows how to overwrite network data with zone association enabled. The first row modifies network "20.0.1.0/24" zone association from "[dnszone1.com](#)" to "[dnszone2.com](#)". The second row modifies network "20.0.2.0/24" default zone association from "[dnszone2.com](#)" to "[dnszone1.com](#)".

```
header-network,address*,netmask*,dhcp_members,network_view,zone_associations
```

```
network,20.0.1.0,255.255.255.0,ib-10-34-43-2.infoblox.com,default,dnszone2.com/True/default
```

```
network,20.0.2.0,255.255.255.0,ib-10-34-43-2.infoblox.com,default,dnszone1.com/True/default, dnszone2.com/False/default
```

### IPv6 Network Container

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-IPv6NetworkContainer	String	Yes			Identifies the first row as a header row for IPv6 network container objects. Example: IPv6NetworkContainer.
address	IP Address	Yes	IP Address	ipv6addr	Indicates the IP address of the network container. Example: 2001::
cidr	Prefix	Yes	Netmask		Indicates the network container, in CIDR format, to which this network container belongs. Example: 64
network_view	String	No	Network View	network_view	If no network view is specified, the default view is used. Example: Default.
comment	String	No	Comment	comment	Example: This is an IPv6 network container.
zone_associations	Zone list	No	DNS Zone Associations	zone_associations	Example: test.com/TRUE

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
valid_lifetime	Integer	No	Valid Lifetime	valid_lifetime	Example: 43200
Preferred_lifetime	Integer	No	Preferred Lifetime	preferred_lifetime	Example: 604800
domain_name	String	No	Domain Name		Example: testdomain.com
domain_name_servers	IP address list	No	DNS Servers		Example: '2000::10,3000::10'
OPTION-7	Integer	No	Custom DHCP Options	override_options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> on page 9. Example: '12' name implies option space = 'DHCPv6', option code/number 7
OPTION-XXXX-200	Option information	No	Custom DHCP Options	override_options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> on page 9. Example: 'dfdfdfd' name implies vendor_class='XXXX', option code/number 200
recycle_leases	Boolean	No	Lease Deletion	recycle_leases	This field is set to TRUE by default. Ensure that you use the overwrite option if you want to change the value to FALSE. Merging data from an import preserves the default value.
enable_ddns	Boolean	No	Enable DDNS Updates	enable_ddns	Example: TRUE
ddns_domainname	String	No	DDNS Domain Name	ddns_domainname	Example: ddns.corp100.com
ddns_ttl	Unsigned integer	No	DDNS Update TTL	ddns_ttl	Indicates the DDNS TTL value in seconds. This is an inherited field. Example: 1200
generate_hostname	Boolean	No	Generate Hostname	override_ddns_generate_hostname	Example: TRUE

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
always_update_dns	Boolean	No	FQDN Support	ddns_server_always_updates	Example: TRUE
update_dns_on_lease_renewal	Boolean	No	Lease Renewal Update	override_update_dns_on_lease_renewal	Example: TRUE
rir_organization	String	No	Organization Name	rir_organization	Identifies the Regional Internet Registry (RIR) organization object. Use this only when the network is associated with an RIR organization. Example: corp100
rir_registration_status	String	No	Registration Status	rir_registration_status	Identifies the registration status of Regional Internet Registry (RIR). Use this only for an RIR network. When you enable the <b>Enable Updates Of RIR Registrations</b> checkbox at the Grid level and import a CSV file to add either an <b>IPv4 network container</b> or an <b>IPv6 network container</b> with the <b>rir_registration_status</b> set to Registered without values for any other RIR fields, the appliance completes the import operation and adds the IPv4 network container or the IPv6 network container to the Grid. The status of this IPv4 network container or the IPv6 network container is set as <b>Non-registered</b> network. Example: Non-registered
last_rir_registration_update_sent	String	No			Identifies the last registration update timestamp of Regional Internet Registry (RIR). This is a read-only attribute.
last_rir_registration_update_status	String	No			Identifies the last registration update status of Regional Internet Registry (RIR). This is a read-only attribute.
enable_discovery	Boolean	Yes	Enable Discovery	network	If this field is set to <b>True</b> , the <b>discovery_member</b> must also be defined.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
discovery_member	String	Yes	Discovery Member	network	This field is required if discovery is enabled for the network.
discovery_exclusion_range	IP Prefix	No	Network Editor → Discovery Exclusions	Network	List of IP ranges to be excluded from the discovery process.
EA-Site	String	No	Extensible attribute	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. Example: California.
EA-Users	String	No	Extensible attribute	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. Example: John.
ADMGRP-XXXX	String	No		permission	ADMGRP-JimSmith is an example of an admin permission of a specific admin group. Example: RW

## IPv6 Network

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-IPv6Network	String	Yes			Example: IPv6Network
rir_organization	String	No	Organization Name	rir_organization	Use this only when the network is associated with an RIR organization. Example: corp100
rir_registration_status	String	No	Registration Status	rir_registration_statuses	Use this only when this is an RIR network. Example: Registered
address	IPv6 address	Yes	Address	ipv6addrs	Example: 2001::
cidr	Prefix	Yes	Netmask		Example: 32

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
network_view	String	No	Network View	network_view	If no view is specified, the Default view is used. Example: Default
enable_discovery	Boolean	Yes	Enable Discovery	enable_discovery	If this field is set to TRUE, the discovery_member must also be defined.
discovery_member	String	Yes	Discovery Member	discovery_member	Required if discovery is enabled for the network.
discovery_exclusion_range	IP Prefix	No	Network Editor → Discovery Exclusions	discovery_exclusion_range	List of IP ranges to be excluded from the discovery process.
disabled	Boolean	No	Disabled	disable	Example: TRUE
comment	String	No	Comment	comment	
auto_create_reversezone	Boolean	No	Automatically create reverse mapping zone	auto_create_reversezone	Example: TRUE
zone_associations	Zone list	No	DNS Zone Associations	zone_associations	Example: <a href="#">test.com/TRUE</a>
dhcp_members	Grid member list	No	Members	members	Example: “ <a href="#">host1.infoblox.com</a> , <a href="#">host2.infoblox.com</a> ”
domain_name	String	No	Domain Name		Example: <a href="#">testdomain.com</a>
domain_name_servers	IP address list	No	DNS Servers		Example: '2000::10,3000::10'
valid_lifetime	Integer	No	Valid Lifetime	valid_lifetime	Example: 43200
Preferred_lifetime	Integer	No	Preferred Lifetime	preferred_lifetime	Example: 604800
recycle_leases	Boolean	No		recycle_leases	Example: FALSE
enable_ddns	Boolean	No	Enable DDNS Updates	enable_ddns	Example: TRUE

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
always_update_dns	Boolean	No	FQDN Support	ddns_server_always_updates	Example: TRUE
ddns_domain_name	String	No	DDNS Domain Name	ddns_domainname	Example: <a href="#">ddnsdomain.com</a>
ddns_ttl	Unsigned integer	No	DDNS Update TTL	ddns_ttl	Example: 3600
generate_hostname	Boolean	No	Generate Hostname	override_ddns_generate_hostname	Example: TRUE
update_dns_on_lease_renewal	Boolean	No	Lease Renewal Update	override_update_dns_on_lease_renewal	Example: TRUE
VLANs	String	No	Assigned VLAN ID Assigned VLAN Name		VLAN View is a container object which can contain VLAN Range/VLAN objects.  Example: default/1/4094/1
OPTION-7	Integer	No	Custom DHCP Options	override_options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> on page 9. Example: '12' name implies option space = 'DHCPv6', option code/number 7
OPTION-XXXX-200	Option information	No	Custom DHCP Options	options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: 'dfdfdf' name implies vendor_class='XXXX', option code/number 200
EA-Site	String	No	Extensible Attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
EA-XXX	String	No	Extensible Attribute XXX value	inheritable	EA-XXX is an example of an inheritable extensible attribute where XXX represents the attribute name. This column displays the value of an inheritable extensible attribute. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EAIherited-XXX	String	No	Inheritance State of an Extensible Attribute XXX	inheritable	EAIherited-XXX is an example of an inheritable extensible attribute where XXX represents the attribute name. This column displays the inheritance state of an inheritable extensible attribute. This column is displayed only if the extensible attribute is inheritable. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Extensible Attribute Users	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for IPv6 networks. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.



## Adding an IPv6 Network

This example shows how to import an IPv6 network 3333::/64 in the Default DHCP view.

```
header-IPv6network,address*,cidr*,network_view,comment,auto_create_reversezone
IPv6network,2001:db8:6000:5000::,64,Default,This is a comment,TRUE
```

For a network that has discovery enabled:

```
header-
IPv6network,address*,cidr*,network_view,auto_create_reversezone,enable_discover
y,discovery_member
IPv6network,2001:db8:6000:5000::,64,Default,TRUE,TRUE,gridmember1.localdomain
```

For a network that has discovery enabled and including an IPv6 range that is excluded for discovery:

```
header-network,address*,netmask*,network_view,auto_create_reversezone,
discovery_enabled,discovery_member,discovery_exclusion_range
IPv6network,2001:db8:6000:5000::,64,Default,TRUE,TRUE,gridmember1.localdomain,
2001:db8:6000:5000::1-2001:db8:6000:5000::128
```

You cannot import a network container, but you can add an IPv6 network container through GUI.

## Overwriting IPv6 Network Data

This example shows how to overwrite the following data of an IPv6 network: address and cidr.

```
header-IPv6network,address*,_new_address,cidr*,_new_cidr
IPv6network,3333::,2222::,64,32
```

## Merging IPv6 Network Data

This example shows how to merge the extensible attribute State and admin group USA\_admins to an IPv6 network.

```
header-IPv6network.address*,cidr*,EA-State,ADMGRP-USA_admins
IPv6network,3333::,64,CA,RW
```

## IPv4 Shared Network

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-SharedNetwork	String	Yes			Example: SharedNetwork
name	String	Yes	Name	name	Example: Site Network
_new_name	String	No			Add this field to overwrite the name field when you select the overwrite or merge option.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
networks	List	Yes			Example: "10.1.1.0/24,10.1.2.0/24," Note that you must first create the IPv4 networks 10.1.1.0/24 and 10.1.2.0/24 before designating them to the shared network.
network_view	String	No	Network View	network_view	If no view is specified, the Default view is used. Example: Default
is_authoritative	Boolean	No	Authoritative	authority	Example: FALSE
option_logic_filters	List of IPv4 logic filter rules		Filter Type/ Action		Examples: .com.infoblox.dns.dhcp_mac_filter\$mac_filter_name, .com.infoblox.dns.nac_filter\$nac_filter_name, .com.infoblox.dns.dhcp_option_filters\$option_filter_name
boot_file	String	No	Boot File	bootfile	Example: bootfile1
boot_server	String	No	Boot Server	bootserver	Example: abc.corp100.com
comment	String	No	Comment	comment	
generate_hostname	Boolean	No	Generate Hostname	ddns_generate_hostname	Example: TRUE
always_update_dns	Boolean	No	DNS Zones Associations	ddns_server_always_updates	Example: FALSE
update_static_leases	Boolean	No	Fixed Address Updates	ddns_update_fixed_address	Example: FALSE
update_dns_on_lease_renewal	Boolean	No	Update DNS on DHCP Lease Renewal	override_update_dns_on_lease_renewal	Example: TRUE
ddns_ttl	Integer	No	DDNS Update TTL	ddns_ttl	This is an inherited field. Example: 1200
enable_option81	Boolean	No	Option 81 Support		ddns_use_option Example: TRUE 81
deny_bootp	Boolean	No	Deny BOOTP Requests	deny_bootp	Example: FALSE
disabled	Boolean	No	Disable	disable	Example: FALSE

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
enable_ddns	Boolean	No	Enable DDNS Updates	enable_ddns	Example: FALSE
ignore_client_requested_options	Boolean	No	Ignore Optionlist	ignore_dhcp_option_list_request	Example: TRUE
next_server	String	No	Next Server	nextserver	Example: blue.domain.com
lease_time	Unsigned integer	No	Lease Time		Example: 1100
enable_pxe_lease_time	Boolean	No	Enable PXE time	lease	Example: FALSE
pxe_lease_time	Unsigned integer	No	PXE Lease Time	pxe_lease_time	Example: 1100
routers	IP address list	No	routers		Example: "10.0.0.1,10.0.0.100"
domain_name	FQDN	No	Domain Name	domain_name	
domain_name_servers	IP address list	No	Name Server		Example: "10.2.3.4,11.2.3.4"
OPTION-2	Integer	No	Custom DHCP Options	options	Example: 50
OPTION-1	String	No	Custom DHCP Options	options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: '255.0.0.0' name implies vendor_class='DHCP' (default)
OPTION-XXXX-200	Option information	No	Extensible Attribute Site	options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: 'dfdfdf' name implies vendor_class='XXXX', option code/number 200
EA-Site	String	No	Extensible Attribute Users	extensible attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Permissions Admin Group/Role	extensible attributes	extensible EA-Users is an example of a user defined

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for shared networks. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding an IPv4 Shared Network

This example shows how to import a new shared Network in the Default DHCP view.

```
header-SharedNetwork,name*,networks*,network_view,comment,OPTION-2
SharedNetwork,Sharednetwork01,"10.0.0.0/24,20.0.0.0/24",Default,This is a
comment.,128
```

### Overwriting IPv4 Shared Network Data

This example shows how to overwrite a shared network with additional networks.

```
header-SharedNetwork,name*,networks*
SharedNetwork,Sharednetwork01,"10.0.0.0/24,20.0.0.0/24,30.0.0.0/24"
```

### Merging IPv4 Shared Network Data

This example shows how to merge the extensible attribute Site and a new admin group DHCP\_admins with RW permission to an IPv4 shared network.

```
header-SharedNetwork,name*,networks*,EA-Site,ADMGRP-DHCP_admins
SharedNetwork,Sharednetwork01,"10.0.0.0/24,20.0.0.0/24",USA,RW
```

## IPv6 Shared Network



### Note

This object is supported in CSV export only.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-IPv6SharedNetwork	String	Yes			Example: IPv6SharedNetwork
name	String	Yes	Name	name	Example: IPv6Shared01

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
_new_name	String	No			Add this field to overwrite the name field when you select the overwrite or merge option.
networks	List	Yes			Example: "2000::/64,3000::/64". Note that you must first create the IPv6 networks 2000::/64 and 3000::/64 before designating them to the shared network.
network_view	String	No	Network View	network_view	If no view is specified, the Default view is used. Example: Default
comment	String	No	Comment	comment	
disabled	Boolean	No	Disabled	disable	Example: FALSE
domain_name	String	No	Domain Name		Example: testdomain.com
domain_name_servers	IP address list	No	Name Server		Example: "2000::10,3000::10"
valid_lifetime	Integer	No	Valid Lifetime	valid_lifetime	Example: 43200
Preferred_lifetime	Integer	No	Preferred Lifetime	preferred_lifetime	Example: 604800
enable_ddns	Boolean	No	Enable DDNS Updates	enable_ddns	Example: TRUE
always_update_dns	Boolean	No	DNS Zones Associations	ddns_server_always_updates	Example: TRUE
ddns_domain_name	String	No	DDNS Domain Name	ddns_domainname	Example: DDNSdomain
ddns_ttl	Integer	No	DDNS Update TTL	override_ddns_ttl	This is an inherited field. Example: 1200
generate_hostname	Boolean	No	Generate Hostname	ddns_generate_hostname	Example: Example: FALSE TRUE
update_dns_on_lease_renewal	Boolean	No	Update DNS on DHCP Lease Renewal	override_update_dns_on_lease_renewal	Example: TRUE

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
OPTION-7	Integer	No	Custom DHCP Options	override_options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: '12' name implies option space = 'DHCPv6', option code/number 7
OPTION-XXXX-200	Option information	No	Custom DHCP Options	override_options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: 'dfdfdf' name implies vendor_class='XXXX', option code/number 200
EA-Site	String	No	Extensible Attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Extensible attribute Users	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## IPv4 DHCP Range

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-DhcpRange	String	Yes			Example: DhcpRange
start_address	IP address	Yes	Start		start_addr Example: 10.0.0.11
_new_start_address	IP address	No			Add this field to overwrite the start_address field when you select the overwrite or merge option. Example: 10.0.0.55
end_address	IP address	Yes	End	end_addr	Example: 10.0.0.22

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
_new_end_address	IP address	No			Add this field to overwrite the end_address field when you select the overwrite or merge option. Example: 10.0.0.66
network_view	String	No	Network View	network_view	If no view is specified, the Default view is used. Example: Default
name	String	No	Name	name	
comment	String	No	Comment	comment	
is_authoritative	Boolean	No	Authoritative	authority	Example: FALSE
boot_file	String	No	Boot File	bootfile	Example: bootfile1
boot_server	String	No	Boot Server	bootserver	Example: abc.corp100.com
ddns_domainname	String	No	DDNS Domain Name	ddns_domainname	Example: ddns.corp100.com
generate_hostname	Boolean	No	Generate Hostname	ddns_generate_hostname	Example: TRUE
deny_all_clients	Boolean	No		deny_all_clients	Example: FALSE
deny_bootp	Boolean	No	Deny BOOTP Requests	deny_bootp	Example: FALSE
disabled	Boolean	No	Disabled	disable	Example: FALSE
domain_name_servers	IP address list	No	Name Servers		Example: "10.2.3.4,11.2.3.4,"
enable_ddns	Boolean	No	Enable DDNS Updates	enable_ddns	Example: FALSE
enable_thresholds	Boolean	No	Enable DHCP Thresholds	enable_dhcp_thresholds	When you set this field to TRUE, you must enter values in the range_high_water_mark and range_low_water_mark fields. You cannot leave those fields empty. Otherwise, the appliance generates an error.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
enable_threshold_email_warnings	Boolean	No	Enable Email Warnings	enable_email_warnings	When you use the merge function, the appliance preserves the existing value in this field. When you use the overwrite function, you must include a value (TRUE or FALSE). Otherwise, the appliance generates an error. Example: TRUE
enable_threshold_snmp_warnings	Boolean	No	Enable SNMP Warnings	enable_snmp_warnings	When you use the merge function, the appliance preserves the existing value in this field. When you use the overwrite function, you must include a value (TRUE or FALSE). Otherwise, the appliance generates an error. Example: TRUE
threshold_email_addresses	email address list	No	Email Addresses		Example: "admin1@infoblox.com',admin2@somewhere.com"
range_high_water_mark	Integer	No	High Water Mark	high_water_mark	When you set enable_thresholds to TRUE, you must enter values in this field and in the range_low_water_mark field. You cannot leave these fields empty. Otherwise, the appliance generates an error. Example: 80
ignore_client_requested_options	Boolean	No	Ignore Optionlist	ignore_dhcp_option_list_request	Example: TRUE
range_low_water_mark	Integer	No	Low Water Mark	low_water_mark	When you set enable_thresholds to TRUE, you must enter values in this field and in the range_high_water_mark field. You cannot leave these fields empty. Otherwise, the appliance generates an error. Example: 10
next_server	String	No	Next Server	nextserver	Example: blue.domain.com
lease_time	Unsigned integer	No	Lease Time		Example: 1100
enable_pxe_lease_time	Boolean	No	Enable PXE lease time		Example: FALSE
pxe_lease_time	Unsigned integer	No	PXE Lease Time	pxe_lease_time	Example: 1100
unknown_clients_option	String	No	Unknown Clients drop-down	unknown_client_option	Example: Allow
known_clients_option	String	No	Known Clients drop-down	known_clients_option	Example: Deny



Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
recycle_leases	Boolean	No	Lease Deletion	recycle_leases	This field is set to TRUE by default. Ensure that you use the overwrite option if you want to change the value to FALSE. Merging data from an import preserves the default value.
update_dns_on_lease_renewal	Boolean	No	Update DNS on DHCP Lease Renewal	override_update_dns_on_lease_renewal	Example: TRUE
always_update_dns	Boolean	No	DNS Zones Associations	always_update_dns	Example: FALSE
exclusion_ranges	IP address range	No	Exclusion Ranges	exclude	This field indicates the start to end address range. You can also include a comment. The valid format is start address-end address/comment. Example: "10.1.0.200-10.1.0.254/"The range for printers', 10.2.3.3-10.2.3.30/"
member	Grid member	No	Served by Grid Member	member	Example: member.infoblox.com
server_association_type	Sting	No			Valid values are MEMBER, NONE, and FAILOVER
failover_association	Sting	No	Served by Failover Association	failover_assoc	
broadcast_address	IP address	No	Broadcast Address		Example: 10.0.0.1
routers	IP address list	No	Routers		Example: "10.0.0.1,10.0.0.100,"
domain_name	FQDN	No	Domain Name	domain_name	
option_logic_filters	List of IPv4 logic filter rules		Filter Type/ Action		Examples: com.infoblox.dns.dhcp_mac_filter\$mac_filter_name, .com.infoblox.dns.nac_filter\$nac_filter_name, .com.infoblox.dns.dhcp_option_filters\$opt_filter_name
OPTION-2	Integer	No	Custom DHCP Options	options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: 50

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
OPTION-1	String	No	Custom DHCP Options	options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: '255.0.0.0' name implies vendor_class='DHCP' (default)
OPTION-XXXX-200	Option information	No	Custom DHCP Options	options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: 'dfdfdfdf' name implies vendor_class='XXXX', option code/number 200
EA-Site	String	No	Extensible Attribute Site	extensible attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-XXX	String	No	Extensible Attribute XXX value	inheritable	EA-XXX is an example of an inheritable extensible attribute where XXX represents the attribute name. This column displays the value of an inheritable extensible attribute. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EAIherited-XXX	String	No	Inheritance State of an Extensible Attribute XXX	inheritable	EAIherited-XXX is an example of an inheritable extensible attribute where XXX represents the attribute name. This column displays the inheritance state of an inheritable extensible attribute. This column is displayed only if the extensible attribute is inheritable. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Extensible Attribute Users	extensible attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for DHCP ranges. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a DHCP Range

This example shows how to import a new DHCP range, Range01, with starting IP of 100.0.0.1 and ending IP 100.0.0.254. The range is assigned to a member, master.corp100.com.

```
header-DhcpRange,start_address*,end_address*,name,comment,member
DhcpRange,100.0.0.1,100.0.0.254,Range01,This is a comment.,master.corp100.com
```

## Overwriting DHCP Range Data

This example shows how to overwrite the starting address and the name of an existing DHCP range.

```
header -DhcpRange, start_address*, _new_start_address, end_address*, name
DhcpRange, 100.0.0.100, 100.0.0.150, 100.0.0.254, Range02
```

## Merging DHCP Range Data

This example shows how to merge an exclusion range 100.0.0.100 to 100.0.0.110 to an existing DHCP range, and to replace a member assignment with a failover association, Failover01.

```
header -
DhcpRange, start_address*, end_address*, exclusion_ranges, failover_association
DhcpRange, 100.0.0.100, 100.0.0.254, 100.0.0.100-100.0.0.110, Fileover01
```

## IPv6 DHCP Range



### Note

This object is supported in CSV export only.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-IPv6DhcpRange	String	Yes			Example: IPv6DhcpRange
address_type	Enumeration	No		address_type	Valid values are ADDRESS, PREFIX, and BOTH. If no value is specified, 'ADDRESS' (default) is used. Example: PREFIX
parent	String	No	Select Network	network	This field is required when address_type is 'PREFIX'. Example: 2000::/16
start_address	IP address	No	Address Start	start_addr	This field is required if address_type is 'ADDRESS' or 'BOTH'. Example: 2000::1
_new_start_address	IP address	No			Add this field to overwrite the start_address field when you select the overwrite or merge option.
end_address	IP address	No	Address End	end_addr	This field is required if address_type is 'ADDRESS' or 'BOTH'. Example: 2000::1

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
_new_end_address	IP address	No			Add this field to overwrite the end_address field when you select the overwrite or merge option.
ipv6_start_prefix	IPv6 address prefix	No	Prefix Delegated Start	ipv6_start_prefix	This field is required if address_type is 'PREFIX' or 'BOTH'. Example: 2000:1111::
_new_ipv6_start_prefix	IPv6 address prefix	No			Add this field to overwrite the ipv6_start_address field when you select the overwrite or merge option.
ipv6_end_prefix	IPv6 address prefix	No	Prefix Delegated End	ipv6_end_prefix	This field is required if address_type is 'PREFIX' or 'BOTH'. Example: 2000:1111::
_new_ipv6_end_prefix	IPv6 address prefix	No			Add this field to overwrite the ipv6_end_address field when you select the overwrite or merge option.
ipb6_prefix_bits	Integer	No		ipv6_prefix_bits	This field is required if address_type is 'PREFIX' or 'BOTH'. Example: 32
network_view	String	No	Network View	network_view	If no view is specified, the Default view is used. Example: Default
name	String	No	Name	name	
comment	String	No	Comment	comment	Example: This is an IPv6 DHCP range.
disabled	Boolean	No	Disabled	disable	Example: FALSE
member	Grid member	No	Grid Member	member	Example: member.infoblox.com
server_association_type	String	No		server_association_type	Valid values are MEMBER and NONE. If no value is specified, None (default) is used.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
exclusion_ranges	IP address range	No	Exclusion Ranges	exclude	This field indicates the start to end address range. You can also include a comment. The valid format is: start address-end address/comment. Example: "2000::1-2000::5/The range for printers"
recycle_leases	Boolean	No	Lease Deletion	recycle_leases	This field is set to TRUE by default. Ensure that you use the overwrite option if you want to change the value to FALSE. Merging data from an import preserves the default value.
EA-Site	String	No	Extensible Attribute Site	extensible attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-XXX	String	No	Extensible Attribute XXX value	inheritable	EA-XXX is an example of an inheritable extensible attribute where XXX represents the attribute name. This column displays the value of an inheritable extensible attribute. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EAIherited-XXX	String	No	Inheritance State of an Extensible Attribute XXX	inheritable	EAIherited-XXX is an example of an inheritable extensible attribute where XXX represents the attribute name. This column displays the inheritance state of an inheritable extensible attribute. This column is displayed only if the extensible attribute is inheritable. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Extensible Attribute Users	extensible attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ADMGRP-JoeSmith	String	No	Permissions Admin Group/ Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## IPv4 Fixed Address/Reservation



### Note

You can use the fixed address header to import reservations. When you import a reservation, you must specify 00:00:00:00:00:00 in the mac\_address field.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-FixedAddress	String	Yes			Example: FixedAddress
ip_address	IP address	Yes	IP Address	ipv4addrss	Example: 10.0.0.11
ms_Server	IP address	Yes			Example: 100.102.30.180
_new_ip_address	IP address	No			Add this field to overwrite the ip_address field when you use the overwrite or merge option.
network_view	String	No	Network View	network_view	If no view is specified, the Default view is used. Example: Default
name	String	No	Name	name	
always_update_dns	Boolean	No	DNS Zones Associations	always_update_dns	Example: FALSE
option_logic_filters	List of IPv4 logic filter rules	No	Filter Type/ Action		Examples: .com.infoblox.dns.dhcp_mac_filter\$mac_filter_name, .com.infoblox.dns.nac_filter\$nac_filter_name, .com.infoblox.dns.dhcp_option_filters\$oapt_filter_name
boot_file	String	No	Boot File	bootfile	Example: bootfile1
boot_server	String	No	Boot Server	bootserver	Example: abc.corp100.com
prepared_zero	Boolean	No			Example: FALSE
comment	String	No	Comment	comment	

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ddns_domainname	String	No	DDNS Domain Name	ddns_domainname	Example: ddns.corp100.com
deny_bootp	Boolean	No	Deny BOOTP Requests	deny_bootp	Example: FALSE
broadcast_address	IP address list	No	Broadcast Address		Example: 10.0.0.1
routers	IP address list	No	Routers		Example: "10.0.0.1,10.0.0.100,"
domain_name	FQDN	No	Domain Name		
domain_name_servers	IP address list	No	Name Servers		Example: "10.2.3.4,11.2.3.4,"
dhcp_client_identifier	String	No	Client Identifier	dhcp_client_identifier	
disabled	Boolean	No	Disabled	disable	Example: FALSE
enable_ddns	Boolean	No	Enable DDNS Updates		Example: FALSE
ignore_client_requested_options	Boolean	No	Ignore Optionlist	ignore_dhcp_option_list_request	Example: TRUE
circuit_id	String	No		agent_circuit_id	This field is required when match_option = CIRCUIT_ID. Example: 11
remote_id	String	No		agent_remote_id	This field is required when match_option = REMOTE_ID. Example: xyz
mac_address	MAC address	No Yes for reservation	MAC Address	mac	This field is required if match_option = MAC_ADDRESS, or if you are importing a reservation. For reservations, you must enter 00:00:00:00:00:00 in this field. Example: aa:bb:cc:dd:ee:ff
match_option	String	No		match_client	Data must be in the following format: ['MAC_ADDRESS','CLIENT_ID','CIRCUIT_ID','REMOTE_ID']
next_server	String	No	Next Server	nextserver	Example: blue.domain.com

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
lease_time	Unsigned integer	No	Lease Time		Example: 1100
enable_pxe_lease_time	Boolean	No	Enable PXE lease time		Example: FALSE
ddns_hostname	String	No		ddns_hostname	Example: host1.test.com
pxe_lease_time	Unsigned integer	No	PXE Lease Time	pxe_lease_time	Example: 1100
OPTION-2	Integer	No	Custom DHCP Options	options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: 50
OPTION-1	String	No	Custom DHCP Options	options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: '255.0.0.0' name implies vendor_class='DHCP' (default)
OPTION-XXXX-200	Option information	No	Custom DHCP Options	options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: 'dfdfdf' name implies vendor_class='XXXX', option code/number 200
EA-Site	String	No	Extensible Attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-XXX	String	No	Extensible Attribute XXX value	inheritable	EA-XXX is an example of an inheritable extensible attribute where XXX represents the attribute name. This column displays the value of an inheritable extensible attribute. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EAIherited-XXX	String	No	Inheritance State of an Extensible Attribute XXX	inheritable	EAIherited-XXX is an example of an inheritable extensible attribute where XXX represents the attribute name. This column displays the inheritance state of an inheritable extensible attribute. This column is displayed only if the extensible attribute is inheritable. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Extensible Attribute User	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .



## Examples

This section contains examples of how to create data files for IPv4 fixed addresses. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding an IPv4 Fixed Address

This example shows how to import a new IPv4 fixed address.

```
header-fixedaddress,ip_address*,network_view,mac_address*,match_option,ms_server,remote_id,EA-Sitefixedaddress,100.0.0.1,Default,aa:aa:aa:aa:aa:aa,Remote_ID,xyz,USA
```

To import MS sync data, specify `ms_server` field in the CSV format.

```
header-fixedaddress,ip_address*,ms_Server,EA-Locationfixedaddress,100.0.0.1,100.102.30.180,location-1
```

### Overwriting IPv4 Fixed Address Data

This example shows how to overwrite the MAC address of an existing IPv4 fixed address from `aa:aa:aa:aa:aa:aa` to `bb:aa:aa:aa:aa:aa`.

```
header-fixedaddress,ip_address*,mac_address*FixedAddress,100.0.0.1,bb:aa:aa:aa:aa:aa
```

This example shows how to overwrite the address of an existing IPv4 fixed address from `100.0.0.1` to `100.0.0.10`.

```
header-fixedaddress,ip_address*,_new_ip_address,mac_address*FixedAddress,100.0.0.1,100.0.0.10,bb:aa:aa:aa:aa:aa
```

### Merging IPv4 Fixed Address Data

This example shows how to merge a new comment to an existing IPv4 fixed address.

```
header-fixedaddress,ip_address,mac_address*,commentFixedAddress,100.0.0.10,bb:aa:aa:aa:aa:aa,A new comment here.
```

This example shows how to merge new routers addresses and the domain name to an existing IPv4 fixed address.

```
header-fixedaddress,ip_address*,mac_address*,routers,domain_namefixedaddress,100.0.0.10,bb:aa:aa:aa:aa:aa,"2.2.2.2,4.4.4.4",ns1.corp100.com
```

### Adding an IPv4 Reservation

This example shows how to import a new IPv4 reservation.

```
header-fixedaddress,ip_address*,network_view,mac_address*,EA-Sitefixedaddress,100.0.0.1,Default,00:00:00:00:00:00,USA
```

## IPv6 Fixed Address

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-IPv6FixedAddress	String	Yes			Example: IPv6FixedAddress
address_type	Enumeration	No		address_type	Valid values are ADDRESS, PREFIX, and BOTH. If no value is specified, 'ADDRESS' (default) is used. Example: PREFIX
parent	String	Yes*	Select Network	network	This field is required when address_type is 'PREFIX'. Example: 2000::/16
ip_address	IP address	Yes*	Address	ipv6addrss	This field is required if address_type is 'ADDRESS' or 'BOTH'. Example: 2000::5
_new_ip_address	IP address	No			Add this field to overwrite the ip_address field when you select the overwrite or merge option.
ipv6_prefix	IPv6 address prefix	Yes*	Prefix Delegated	ipv6prefix	This field is required if address_type is 'PREFIX' or 'BOTH'. Example: 2000:1111::
_new_ipv6_prefix	IPv6 address prefix	No			Add this field to overwrite the ipv6_prefix field when you select the overwrite or merge option.
ipv6_prefix_bits	Integer	No		ipv6_prefix_bits	This field is required if address_type is 'PREFIX' or 'BOTH'. Example: 32
network_view	String	No	Network View	network_view	If no view is specified, the Default view is used. Example: Default
name	String	No	Name	name	Example: IPv6FixedAddr
comment	String	No	Comment	comment	
disabled	Boolean	No	Disabled	disable	Example: FALSE
match_option	String	No		match_client	Only 'DUID' is allowed. Example: DUID
duid	String	Yes	DUID	duid	Example: 0001

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
domain_name	FQDN	No	Domain Name		Example: ns1.corp100.com
domain_name_servers	IPv6 address list	No	Name Servers		Example: 2000::10,3000::10
valid_lifetime	Unsigned integer	No	Valid Lifetime	valid_lifename	Example: 43200
preferred_lifetime	Unsigned integer	No	Preferred Lifename	preferred_lifetime	Example: 604800
OPTION-7	Integer	No	Custom DHCP Options	override_options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: '12' name implies option space = 'DHCPv6', option code/number 7
OPTION-XXXX-200	Option information	No	Custom DHCP Options	override_options	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: 'dffd' name implies vendor_class='XXXX', option code/number 200
EA-Site	String	No	Extensible Attribute Site	extensible attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-XXX	String	No	Extensible Attribute XXX value	inheritable	EA-XXX is an example of an inheritable extensible attribute where XXX represents the attribute name. This column displays the value of an inheritable extensible attribute. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EAIherited-XXX	String	No	Inheritance State of an Extensible Attribute XXX	inheritable	EAIherited-XXX is an example of an inheritable extensible attribute where XXX represents the attribute name. This column displays the inheritance state of an inheritable extensible attribute. This column is displayed only if the extensible attribute is inheritable. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
EA-Users	String	No	Extensible Attribute Users	extensible attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/ Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .



#### Note

\* Required in some cases, see detailed field description.

## Examples

This section contains examples of how to create data files for IPv6 fixed addresses. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding an IPv6 Fixed Address

This example shows how to import a new IPv6 fixed address.

```
header-IPv6FixedAddress,address_type,ip_address,network_view,match_option,EA-
Site IPv6FixedAddress,ADDRESS,2000::5,default,DUID,USA
```

### Overwriting IPv6 Fixed Address Data

This example shows how to overwrite an existing IPv6 fixed address from 2000::5 to 2000::1.

```
header-IPv6FixedAddress,address_type,ip_address,_new_ip_address
IPv6FixedAddress,ADDRESS,2000::5,2000::1
```

### Merging IPv6 Fixed Address Data

This example shows how to merge a new comment to an existing IPv6 fixed address.

```
header-IPv6FixedAddress,ip_address,comment IPv6FixedAddress,2000::1,A new
comment.
```

This example shows how to merge a new domain name to an existing IPv6 fixed address.

```
header-IPv6FixedAddress,ip_address,domain_name
IPv6FixedAddress,2000::1,ns1.corp100.com
```

## DHCP Fingerprint

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-DHCPFingerprint	String	Yes			Example: DHCPFingerprint
name	String	Yes	Name	name	Example: Samsung Android
new_name	String	No			Add this field to overwrite the name field when you select the overwrite or merge option.
type	String	No			Device type. This can be 'Standard' or 'Custom.' The default is 'Custom.' Example: Custom
comment	String	No	Comment	comment	
disable	Boolean	No	Disabled	disable	Example: FALSE
vendor_id	String	No	Vendor ID	vendor_id	Example: MSFT 7.x
option_sequence	String	No	Option Number Sequence	option_sequence	DHCP options from 1 to 255 separated by commas (without spaces). Also enter the protocol (ipv4 or ipv6) at the end. Example: "[1,3,6,7,12,15,28,40,41,42,225,226,227,22/ipv4]"
device_class	String	No	Device Class	device_class	Device category to which the DHCP fingerprint belongs. This is used for filtering purposes. Example: Printers
protocol	String	Yes	protocol	protocol	Protocol type. This can be IPV4 or IPV6.
EA-Site	String	No	Extensible Attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to custom DHCP fingerprints. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

### Examples

This section contains examples of how to create data files for DHCP fingerprints. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a Custom DHCP Fingerprint

This example shows how to import a new custom DHCP fingerprint.

```
header-DHCPFingerprint,name*,protocol*,comment,disable,option_sequence
DHCPFingerprint,SpecialPrinter,IPV4,Special printer for
QA,FALSE,['1,3,6,7,12,15,28,40,41,42,225,226,227,228/ipv4']"
```

### Overwriting DHCP Fingerprint Data

This example shows how to overwrite an existing custom DHCP fingerprint.

```
header-DHCPFingerprint,name*,new_name,comment
DHCPFingerprint,SpecialPrinter,QAPrinter,Changed from Special Printer to QA
Printer
```

### Merging DHCP Fingerprint Data

This example shows how to merge a new comment to an existing DHCP fingerprint.

```
header-DHCPFingerprint,name*,comment DHCPFingerprint,Xbox,Blocked from network
```

## DHCP MAC Filter

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-DhcpMacFilter	String	Yes			Example: DhcpMacFilter
name	String	Yes	Name	name	Example: MAC filter 1
_new_name	String	No			Add this field to overwrite the name field when you select the overwrite or merge option.
never_expires	Boolean	No	Never Expires		Example: TRUE
expiration_interval	Integer	No	Automatically expires in	default_mac_address_expiration	Example: 3624
enforce_expiration_time	Boolean	No	Enforce Expiration Times	enforce_expiration_times	Example: FALSE
comment	String	No	Comment	comment	

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
EA-Site	String	No	Extensible Attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Extensible Attribute Users	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for DHCP MAC filters. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a DHCP MAC Filter

This example shows how to import a DHCP MAC filter.

header-

```
dhcpmacfilter,name*,never_expires,expiration_interval,enforce_expiration_time,comment
```

```
dhcpmacfilter,MacFilter01,FALSE,3600,TRUE,This is a comment.
```

### Overwriting DHCP MAC Filter Data

This example shows how to overwrite the MAC filter name from MacFilter01 to MacFilter02.

```
header-dhcpmacfilter,name*,_new_name dhcpmacfilter,MacFilter01,MacFilter02
```

### Merging DHCP MAC Filter Data

This example shows how to merge extensible attributes Site and Users, as well as admin group DHCP\_admins with a RO permission.

```
header-dhcpmacfilter,name*,EA-Site,EA-Users,ADMGRPDHCP_admins
```

```
dhcpmacfilter,MacFilter02,USA,John Smith,RO
```

## MAC Filter Address

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-MacFilterAddress	String	Yes			Example: MacFilterAddress
parent	String	Yes	MAC address filter	filter_name	Example: MAC filter 1
mac_address	MAC address	Yes	MAC address	mac_address	Example: aa:bb:cc:dd:ee:ff
_new_mac_address	MAC address	No			Add this field to overwrite the mac_address field when you select the overwrite or merge option.
is_registered_user	Boolean	No	Register as User		Example: TRUE
registered_user	String	No	Register as User		Example: John Doe
guest_first_name	String	No	Register as Guest: First Name		Example: John
guest_middle_name	String	No	Register as Guest: Middle Name		Example: Doe
guest_last_name	String	No	Register as Guest: Last Name		Example: Doe
guest_email	Email address	No	Register as Guest: Email Address		Example: jdoe@infoblox.com
guest_phone	String	No	Register as Guest: Phone Number		Example: 408-111-1111
guest_custom_field1	String	No	Register as Guest: Custom Field 1		
guest_custom_field2	String	No	Register as Guest: Custom Field 2		
guest_custom_field3	String	No	Register as Guest: Custom Field 3		



Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
guest_custom_field4	String	No	Register as Guest: Custom Field 4		
never_expires	Boolean	No	Never Expires		Example: FALSE
expire_time	Date/Time	No	Expires On		Data must be in the following format: CCYY-MM-DDThh:mm:ss Example: 2009-02-29T10:30:00 The timestamp must be based on UTC time.
comment	String	No	Comment	comment	
EA-Site	String	No	Extensible Attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Extensible Attribute Users	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	permission	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for DHCP MAC filters. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a DHCP MAC Filter

This example shows how to import MAC filter address, aa:aa:aa:aa:aa:aa, to MacFilter01.

header-

```
macfilteraddress,parent*,mac_address*,is_registered_user,never_expires,comment
macfilteraddress,MacFilter01,aa:aa:aa:aa:aa:aa,FALSE,TRUE,This is a comment.
```

## Overwriting DHCP MAC Filter

This example shows how to overwrite a MAC filter address with an expiration time.

```
header-macfilteraddress,parent*,mac_address*,never_expires,expire_time
macfilteraddress,MacFilter01,aa:aa:aa:aa:aa:aa,FALSE,2010-12-30T10:30:00Z
```

## Merging DHCP MAC Filter

This example shows how to merge extensible attributes Site and Users to an existing MAC filter address.

```
header-macfilteraddress,parent*mac_address*,EA-Site,EA-Users
macfilteraddress,MacFilter01,aa:aa:aa:aa:aa:aa,USA,John Smith
```

## Option Filter

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-OptionFilter	String	Yes			Example: OptionFilter
name	String	Yes	Name	name	Example: Option Filter 1
_new_name	String	No			Add this field to overwrite the name field when you select the overwrite or merge option.
comment	String	No	Comment	comment	
boot_file	String	No	Boot File	boot_file	Example: bootfile1
boot_server	String	No	Boot Server	boot_server	Example: abc.corp100.com
lease_time	Integer	No	Lease Time	lease_time	Example: 7200
pxe_lease_time	Unsigned integer	No	PXE Lease Time	pxe_lease_time	Example: 1100
next_server	String	No	Next Server	next_server	Example: blue.domain.com
option_space	String	No	Option Space	option_space	Example: Infoblox_DHCP
OPTION-2	Integer	No	Custom DHCP Options	option_list	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: 50

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
OPTION-1	String	No	Custom DHCP Options	options_list	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: '255.0.0.0' name implies vendor_class='DHCP' (default)
OPTION-XXXX-200	Option information	No	Extensible Attribute Site	options_list	This is an example of a DHCP option. For information, see <a href="#">Data Specific Guidelines</a> . Example: 'dfdfdf' name implies vendor_class='XXXX', option code/number 200
EA-Site	String	No	Extensible Attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Extensible Attribute Users	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for option filters. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding an Option Filter

This example shows how to import an option filter with boot file, boot server, and lease time.

```
header-OptionFilter name*,comment,boot_file,boot_server,lease_time
OptionFilter,OptionFilter01,This is a comment.,bootfile01 1.2.3.4,12800
```

### Overwriting Option Filter Datae

This example shows how to overwrite an option filter name and boot file name.

```
header-OptionFilter,name*,_new_name,boot_file
OptionFilter,OptionFilter01,OptionFilter02,bootfile02
```

## Merging Option Filter Data

This example shows how to merge to an option filter the PXE lease time and next server domain name.

```
header-OptionFilter,name*,pxe_lease_time,next_server
OptionFilter,OptionFilter02,12800,next.corp100.com
```

## Option Filter Match Rule

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-OptionFilterMatchRule	String	Yes			Example: OptionFilterMatchRule
parent	Option filter	Yes	Option Filter Name	filter	The name of the parent option filter. Example: Option filter 1
match_option	String	Yes	Match Option	num	Example: OPTION-1 (option space is DHCP)
match_value	String	Yes	Match Value	value	Example: 255.0.0.0
_new_match_value	String	No			Add this field to overwrite the match_value field when you select the overwrite or merge option.
comment	String	No	Comment	comment	
is_substring	Boolean	No	Value is a substring	is_substring	Example: FALSE
substring_offset	Integer	No	Substring Offset	substring_offset	Example: 2
substring_length	Unsigned integer	No	Substring Length	substring_length	Example: Doe

## Examples

This section contains examples of how to create data files for option filter match rules. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding an Option Filter Match Rule

This example shows how to import an option filter match rule.

```
header-
OptionFilterMatchRule,parent*,match_option*,match_value*,is_substring,substring
_offset
OptionFilterMatchRule,OptionFilter01,OPTION-1,2.2.2.2,FALSE,0
```

### Overwriting Option Filter Match Rule Data

This example shows how to overwrite an existing match option with OPTION-2 and add a new match value of 3.3.3.3.  
header-OptionFilterMatchRule,parent\*,match\_option\*,match\_value\*,\_new\_match\_value

```
OptionFilterMatchRule,OptionFilter01,OPTION-2,2.2.2.2,3.3.3.3
```

### Merging Option Filter Match Rule Data

This example shows how to merge a substring length to an existing option filter match rule.

```
header-
```

```
OptionFilterMatchRule,parent*,match_option*,match_value*,substring_length
```

```
OptionFilterMatchRule,OptionFilter02,OPTION-2,3.3.3.3,256
```

## Host Record

If only one IPv4 address is specified in the host record, you can add DHCP options to the host address.



### Note

IDN is supported for object types: fqdn and aliases. You can use punycode or IDNs while importing these objects.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines Method
Header-HostRecord	String	Yes			Example: HostRecord
fqdn	FQDN	Yes	Name	name zone	This field combines the host record name and the zone name to form the FQDN. Example: h2.corp100.com
_new_fqdn	FQDN	No			Add this field to overwrite the fqdn field when you select the overwrite or merge option.
view	String	No	DNS View	views	If no view is specified, the Default view is used. Example: Default
network_view	String	No	Network View		If no network view is specified, the Default view is used. Example: Default

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines Method
addresses	IP address or IP address list	No	IP Addresses	ipv4addr	You must specify a value in this field or in the ipv6_addresses field. Example: '1.2.3.4' or "1.2.3.4, 5.6.7.8". If there are multiple addresses in the host record, use HostAddress to specify the parameters for each address. For information, see <a href="#">IPv4 Host Address</a> and <a href="#">IPv6 Host Address</a> .
ipv6_addresses	IP address or IP address list	No	IP Addresses	ipv6addr	You must specify a value in this field or in the addresses field. If there are multiple addresses in the host record, use HostAddress to specify the parameters for each address. For information, see <a href="#">IPv4 Host Address</a> and <a href="#">IPv6 Host Address</a> .
aliases	Alias list	No	Aliases	aliases	Example: www.infoblox.com
configure_for_dns	Boolean	No	Enable in DNS	configure_for_dns	Example: TRUE
_new_configure_for_dns	Boolean	No	Enable in DNS	configure_for_dns	Add this field to overwrite the configure_for_dns field when you select the overwrite or merge option.
comment	String	No	Comment	comment	
disabled	Boolean	No	Disable	disabled	Example: FALSE
ttl	Unsigned integer	No	TTL	ttl	This is an inherited field. For information, see <a href="#">Data Specific Guidelines</a> . Example: 28800
mac_address	MAC address	No	MAC Address	mac_address	This field applies to the host address. This is required if the IP address is configured for DHCP. Example: aA:Bb:c2:DD:E1:FF
ddns_protected	Boolean	No	Protected	ddns_protected	If the record is marked as protected, DDNS updates are restricted to only the record. This applies to both static and dynamic records.
configure_for_dhcp	Boolean	No	DHCP checkbox	configure_for_dhcp	This field applies to the host address. Example: TRUE
deny_bootp	Boolean	No	Deny BOOTP Requests	deny_bootp	This field applies to the host address. Example: FALSE
broadcast_address	String	No	Broadcast Address		This field applies to the host address.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines Method
boot_file	String	No	Boot File	boot_file	This field applies to the host address.
boot_server	String	No	Boot Server	boot_server	This field applies to the host address.
next_server	String	No	Next Server	next_server	This field applies to the host address.
lease_time	Unsigned integer	No	Lease Time	lease_time	This field applies to the host address.
pxe_lease_time_enabled	Boolean	No	Enable PXE Lease Time		This field applies to the host address.
pxe_lease_time	Unsigned integer	No	PXE Lease Time	pxe_lease_time	This field applies to the host address.
domain_name	FQDN	No	Domain Name	domain_name	This field applies to the host address.
domain_name_servers	IP list	No	Name Server		This field applies to the host address. Example: "5.6.7.8,1.2.3.4"
routers	IP list	No	Router		This field applies to the host address. Example: "2.0.0.2,1.2.3.4"
match_option	String	No		match_client	This field applies to the host address. Data must be in the following format: MAC_ADDRESS/RESERVED
ignore_dhcp_param_request_list	Boolean	No	Ignore Optionalist	ignore-dhcp_option_list_request	This field applies to the host address.
OPTION-1	String	No	Custom DHCP Options	options	This field applies to the host address. Example: '255.0.0.0' name implies vendor_class='DHCP' (default )
OPTION-XXXX-200	Option Information	No	Custom DHCP Options	options	This field applies to the host address. Example: 'dfdfdfd' name implies vendor_class='XXXX', optioncode/number 200
EA-Site	String	No	Extensible attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines Method
ADMGRP-JoeSmith	String	No	Permissions Admin Group/Role	Permissions Admin Group/Role	ADMGRP-JoeSmith is an example of an admin permission of a specific admin group. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for host records. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a Host Record

This example shows how to add a new host record, host1, in zone corp100.com and DNS view Initial.

```
header-hostrecord,addresses,configure_for_dns*,fqdn*,view
hostrecord,100.0.0.101,TRUE,host1.corp100.com,Initial
```

This example shows how to add a new host record, host2, in zone corp100.com with aliases = www.corp100.com, comment = USA, and TTL = 3600 seconds.

```
header-hostrecord,addresses,configure_for_dns*,fqdn*,aliases,comment,ttl
hostrecord,100.0.0.102,TRUE,host2.corp100.com,www.corp100.com,USA,3600
```

### Overwriting Host Record Data

This example shows how to overwrite the FQDN of an existing host record from host1.corp100.com to new\_host1.corp100.com, and to change the TTL to 128 seconds.

```
header-hostrecord,addresses,configure_for_dns*,fqdn*,_new_fqdn,ttl
hostrecord,100.0.0.101,TRUE,host1.corp100.com,new_host1.corp100.com,1280
```

This example shows how to overwrite the aliases of a host record from www.corp100.com to www.corp200.com and comment from USA to Japan.

```
header-hostrecord,addresses*,configure_for_dns*,fqdn*,aliases,comment
hostrecord,100.0.0.102,TRUE,host2.corp100.com,www.corp200.com,Japan
```

Note that overwriting the host record data is not supported for non-DNS hosts.

### Merging Host Record Data

This example shows how to disable an existing host record.

```
header-hostrecord,configure_for_dns*,addresses*,fqdn*,disabled
hostrecord,100.0.0.101,TRUE,new_host1.corp100.com,TRUE
```

This example shows how to add additional aliases to a host record.

```
header-hostrecord,configure_for_dns*,addresses*,fqdn*,aliases
hostrecord,100.0.0.102,TRUE,host2.corp100.com,"www.corp200.com,http.corp200.com"
```



## Relay Agent Filter

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-RelayAgentFilter	String	Yes			Example: RelayAgentFilter
name	String	Yes	Name	name	Example: Relay Agent Filter 1
_new_name	String	No			Add this field to overwrite the name field when you select the overwrite or merge option.
comment	String	No	Comment	comment	
circuit_id_rule	String	No	Circuit ID: Match Value	is_circuit_id	Example: MATCHES_VALUE
circuit_id	String	No	Circuit ID	circuit_id_name	
remote_id_rule	String	No	Remote ID: Match Value	is_remote_id	Example: MATCHES_VALUE
remote_id	Integer	No	Remote ID	remote_id_name	Example: 50
EA-Site	String	No	Extensible Attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Extensible Attribute Users	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

### Examples

This section contains examples of how to create data files for relay agent filters. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

#### Adding a Relay Agent Filter

This example shows how to import a relay agent filter with a circuit ID and a remote ID.

header-

```
RelayAgentFilter,name*,circuit_id_rule,circuit_id,remote_id_rule,remote_id
```

```
RelayAgentFilter,relayagent1,MATCHES_VALUE,123456,MATCHES_VALUE,abcd
```

## Overwriting Relay Agent Filter Data

This example shows how to overwrite the circuit ID and remote ID of an existing relay agent filter.

```
header-RelayAgentFilter,name*,circuit_id_rule,circuit_id,remote_id_rule
remote_id RelayAgentFilter,relayagent1,MATCHES_VALUE,336699,MATCHES_VALUE,xyz
```

## Merging Relay Agent Filter Data

This example shows how to merge a comment and extensible attribute Site to an existing relay agent filter.

```
header-RelayAgentFilter,name*,comment,EA-Site RelayAgentFilter,relayagent1,This
is a comment.,USA
```

## DHCP Fingerprint Filter

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-DHCPFingerprintFilter	String	Yes			Example: DHCPFingerprintFilter
name	String	Yes	Name	name	Example: HP Printers
_new_name	String	No			Add this field to overwrite the name field when you select the overwrite or merge option.
fingerprint	String	Yes	DHCP Fingerprints	fingerprint	An array of DHCP fingerprints
_new_fingerprint	String	No			Add this field to overwrite the fingerprint field when you select the overwrite or merge option.
comment	String	No	Comment	comment	
EA-Site	String	No	Extensible Attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for DHCP fingerprint filters. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a DHCP Fingerprint Filter

This example shows how to import a DHCP fingerprint filter.

```
header-DHCPFingerprintFilter,name*,fingerprint*,comment
```

```
DHCPFingerprintFilter,WindowsXP1,MSFT5.x,MSFT,Some Windows XP systems
```

### Overwriting DHCP Fingerprint Filter Data

This example shows how to overwrite the filter name of an DHCP fingerprint filter.

```
header-DHCPFingerprintFilter,name*,_new_name
```

```
DHCPFingerprintFilter,WindowsXP1,WindowsXP5
```

### Merging DHCP Fingerprint Filter Data

This example shows how to merge a comment to an DHCP fingerprint filter.

```
header-DHCPFingerprintFilter,name*,comment,EA-Site
```

```
DHCPFingerprintFilter,WindowsXP1,This is a comment
```

## NAC Filter

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-NACFilter	String	Yes			Example: NACFilter
name	String	Yes	Name	name	Example: NAC Filter
_new_name	String	No			Add this field to overwrite the name field when you select the overwrite or merge option.
comment	String	No	Comment	comment	
expression	Option list	No	Match the following rule	expression	To include the option list in the <b>Matching the following rules</b> table in Grid Manager, you must enclose the expression in brackets. Example: (Sophos.ComplianceState="Compliant" AND RADIUS.ServerError="TRUE")
EA-Site	String	No	Extensible Attribute Site	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .
EA-Users	String	No	Extensible Attribute Users	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. For information about data format and examples, see <a href="#">Data Specific Guidelines</a> .

## Examples

This section contains examples of how to create data files for NAC filters. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a NAC Filter

This example shows how to import a new NAC filter.

```
header-NacFilter,name*,comment,expression,EA-Site NacFilter,nacfilter01,This is  
a comment.,option ServerError="true",USA
```

### Overwriting NAC Filter Data

This example shows how to overwrite the name and comment of an existing NAC filter.

```
header-NacFilter,name*,_new_name,comment NacFilter,nacfilter01,nacfilter02,This  
is a new comment.
```

### Merging NAC Filter Data

This example shows how to merge the extensible attribute Users to an existing NAC filter.

```
header-NacFilter,name*,EA-Users NacFilter,nacfilter02,John Smith
```

## IPv4 Option Space

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-OptionSpace	String	Yes			Example: OptionSpace
name	String	Yes	Name	name	Example: ABC-co options
_new_name	String	No			Add this field to overwrite the name field when you select the overwrite or merge option.
comment	String	No	Comment	comment	

### Examples

This section contains examples of how to create data files for option spaces. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

#### Adding an IPv4 Option Space

This example shows how to import a new IPv4 option space.

```
header-OptionSpace,name*,space_type OptionSpace,Optionspace01,VENDOR_SPACE
```

#### Overwriting IPv4 Option Space Data

This example shows how to overwrite the IPv4 option space name with Optionspace02.

```
header-OptionSpace,name*,_new_name OptionSpace,Optionspace01,Optionspace02
```

### Merging IPv4 Option Space Data

This example shows how to merge a comment to the IPv4 option space Optionspace02.

```
header-OptionSpace,name*,comment OptionSpace,Optionspace02,This is a comment.
```

## IPv6 Option Space

### Note

This object is supported in CSV export only.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-IPv6OptionSpace	String	Yes			Example: IPv6OptionSpace
name	String	Yes	Name	name	Example: MySpace
_new_name	String	No			Add this field to overwrite the name field when you select the overwrite or merge option.
comment	String	No	Comment	comment	Example: Custom option space
ipv6_enterprise_number	String	No	Enterprise Number		This is the vendor's enterprise number that is registered with IANA. Example: 7779

## IPv4 Option Definition

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-OptionDefinition	String	Yes			Example: OptionDefinition
space	String	Yes	Option Space	space	Example: ABC-co options
_new_space	String	No			Add this field to overwrite the space field when you select the overwrite or merge option.
name	String	Yes	Name	name	Example: Option one
_new_name	String	No			Add this field to overwrite the name field when you select the overwrite or merge option.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
code	String	Yes	Options: Code	code	
type	String	Yes	Options: Type	type	Use any of the following values: T_FLAG, T_STRING, T_TEXT, T_IP_ADDRESS, T_ARRAY_IP_ADDRESS, T_DOMAIN, T_ARRAY_DOMAIN, T_UINT8, T_UINT16, T_UINT32, T_INT8, T_INT16 Example: T_TEXT

## Examples

This section contains examples of how to create data files for option definitions. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding an IPv4 Option Definition

This example shows how to import an option definition to the IPv4 option space Space01.

```
header-OptionDefinition,space*,name*,code*,type*
OptionDefinition,Space01,Option99,99,T_IP_ADDRESS
```

### Overwriting IPv4 Option Definition Data

This example shows how to overwrite the IPv4 option definition type from T\_IP\_ADDRESS to T\_FLAG.

```
header-OptionDefinition,space*,name*,code*,type*
OptionDefinition,Space01,Option99,99,T_FLAG
```

### Merging IPv4 Option Definition Data

This example shows how to merge a new space and a new name to an existing IPv4 option definition.

```
header-OptionDefinition,space*,_new_space,name*,_new_name
OptionDefinition,Space01,New_Space01,Option99,New_Option99
```

## IPv6 Option Definition

### Note

This object is supported in CSV export only.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-IPv6OptionDefinition	String	Yes			Example: IPv6OptionDefinition

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
space	String	Yes	Option Space	space	Example: MySpace
_new_space	String	No			Add this field to overwrite the space field when you select the overwrite or merge option.
name	String	Yes	Name	name	Example: MyOption1
_new_name	String	No			Add this field to overwrite the name field when you select the overwrite or merge option.
code	String	Yes	Options: Code	code	Example: 10
type	String	Yes	Options: Type	type	Use any of the following values: 'T_ARRAY_DOMAIN', 'T_ARRAY_INT16', 'T_ARRAY_INT32', 'T_ARRAY_INT8', 'T_ARRAY_IP_ADDRESS', 'T_ARRAY_IP_ADDRESS_PAIR', 'T_ARRAY_UINT16', 'T_ARRAY_UINT32', 'T_ARRAY_UINT8', 'T_DOMAIN', 'T_FLAG', 'T_FLAG_IP_ADDRESS', 'T_FLAG_TEXT', 'T_INT16', 'T_INT32', 'T_INT8', 'T_IP_ADDRESS', 'T_STRING', 'T_TEXT', 'T_UINT16', 'T_UINT32', 'T_UINT8', 'T_UINT8_1_2_4_8' Example: T_INT8

## Permissions for DNS Resources with Associated IP Addresses in Networks and Ranges

You can further control permissions for DNS resources that have associated IP addresses in a network container, network, or address range. These DNS resources include A records, AAAA records, PTR records, and DNS hosts. Permissions for these resources have been added so you now have more control over who can perform which tasks for these DNS resources without affecting permissions defined for the networks and ranges to which the resources belong. For more information about this feature, refer to the *Infoblox Administrator Guide*.

As a superuser, you can now grant permissions to admin groups for more granular access to the following resources:

- IPv4 and IPv6 DHCP fixed addresses and IPv4 reservations in a range
  - IPv4 and IPv6 host addresses in a range
  - A and AAAA records in a network container, network, or range
  - IPv4 and IPv6 PTR records in a network container, network, or range
- Following are some examples:

Permissions for Host Address and Fixed Address in a DHCP Network:

```
header-network, address*, netmask*, ADMGRP-foogroup
network, 10.100.0.0, 255.255.0.0, "RW, HostAddress/RW, FixedAddress/DENY"
header-IPv6network, address*, cidr*, ADMGRP-foogroup
IPv6network, 2001::, 64, "RW, IPv6FixedAddress/RW, IPv6HostAddress/DENY"
```

Permissions for Host Address and Fixed Address in a DHCP Range:

```
header-dhcprange, start_address*, end_address*, ADMGRP-foogroup
dhcprange, 10.100.20.0, 10.100.20.255, "DENY, FixedAddress/RO, HostAddress/RW"
```

Permissions for A and PTR Records in a DHCP Network:

```
header-network, address*, netmask*, ADMGRP-foogroup
network, 30.30.0.0, 255.255.0.0, "RW, ARecord/RW, PtrRecord/DENY"
```



**Note**

You cannot import network containers and IPv6 ranges using CSV import.

## DHCP Failover Association

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-DhcpFailoverAssoc	String	Yes			Identifies the first row as a header row for the DHCP failover association. Example: DhcpFailoverAssoc.
name	String	Yes	Name	name	Indicates the name of the DHCP failover association. Example: dhcp_fo_assoc1.
_new_name	String	No			Add this field to overwrite the <b>name</b> field when you select the <b>Overwrite</b> or <b>Merge</b> option.
comment	String	No	Comment	comment	Example: DHCP Failover Association.
primary_server_type	String	Yes			Indicates whether the primary server of the name server group is set to Grid or External. Example: GRID
grid_primary	String	No	Grid Primary	primary	Indicates the name of the Grid primary. Example: infoblox.localdomain
external_primary	IP address	No	External Primary	primary	Indicates the IP address. Example: 10.10.10.1
secondary_server_type	String	Yes			Indicates whether the secondary server of the name server group is set to Grid or External. Example: External
grid_secondary	String	No	Grid Secondary	secondary	Indicates the name of the Grid secondary. Example: infoblox2.localdomain



Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
external_secondary	IP address	No	External Secondary	secondary	Indicates the IP address. Example: 20.20.20.1
failover_port	Unsigned integer	No			Indicates the failover port number. The default value is 647. The port number must be between 1 and 63999.
max_response_delay	Unsigned integer	No		max_response_delay	Indicates the maximum response delay value of a DHCP failover object. Default value is 60 seconds and the minimum value is one second.
max_unacked_updates	Unsigned integer	No		max_unacked_updates	Indicates the maximum number of unacked updates value of a DHCP failover object. Default value is 10 minutes and the minimum value is one second.
mclt	Unsigned integer	No		max_client_lead_time	Indicates the maximum client lead time value of a DHCP failover object. Specify the value of the maximum client lead time in a 32-bit integer format (range from 0 to 4294967295) that represents the duration in seconds. The default value is 3600.
max_load_balance_delay	Unsigned integer	No		max_load_balance_delay	Indicates the maximum load balancing delay value of a DHCP failover object. Specify the value of the maximum load balancing delay in a 32-bit integer format (range from 0 to 4294967295) that represents the duration in seconds. The default value is three seconds.
load_balance_split	Unsigned integer	No		load_balance_split	Indicates the load balancing split value of a DHCP failover object. Specify a value from zero to 255. The default value is 128.
recycle_leases	Boolean	No	Lease Deletion	recycle_leases	This field is set to <b>True</b> by default. When you set this to <b>True</b> , leases in a deleted range are kept until expiration. Ensure that you use the <b>Overwrite</b> option if you want to change the value to <b>False</b> . Merging data from an import preserves the default value.
EA-Site	String	No	Extensible attribute	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. Example: California.
EA-Users	String	No	Extensible attribute	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. Example: John.

## Grid Member

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-Member	String	Yes			Identifies the first row as a header row for the member objects. Example: Member

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
_new_name	FQDN	No			Add this field to overwrite the <b>name</b> field when you select the <b>Overwrite</b> or <b>Merge</b> option.
comment	String	No	Comment	comment	Example: This is a Grid member.
enable_ha	Boolean	No		enable_ha	Enable or disable HA. This is <b>True</b> if HA is enabled. Example: False
ha_nodes	String	No			Indicates the list of ha_status, or public_ip_address, or ipv6_public_ip_address, or ha_ip_address. Example: ACTIVE/10.0.0.11//10.0.0.12, PASSIVE/10.0.0.13//10.0.0.14
vpn_mtu	Unsigned integer	No		vpn_mtu	Indicates maximum transmission unit of the VPN. Example: 1450
ipv4addr	String	Yes (Insert)	IP Address	ipv4addr	Indicates IPv4 address. Example: 10.0.0.10. Note that you must specify this column in the CSV file when you perform a CSV Import Insert operation.
mask	String	Yes (Insert)		mask	Indicates the netmask. Example: 255.255.255.0. Note that you must specify this column in the CSV file when you perform a CSV Import Insert operation.
gateway	String	Yes (Insert)		gateway	Indicates the gateway address. Example: 10.0.0.1. Note that you must specify this column in the CSV file when you perform a CSV Import Insert operation.
vlan_id	Unsigned integer	No			VLAN Id for LAN1 and LAN2 address. Example: 10
ipv6addr	String	No	IP Address	ipv6addr	Indicates IPv6 address. Example: 2001::10
ipv6_cidr	Unsigned integer	No		ipv6_cidr	The CIDR of the Grid member. This is required only when ipv6addr is specified. Example: 64
ipv6_enable_auto_config	Boolean	No		ipv6_enable_auto_config	Enable or disable IPv6 auto-configuration of the Grid member. Example: False

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ipv6_gateway	String	No		ipv6_gateway	The LAN IPv6 gateway of the Grid member. Example: 2001::1
ipv6_vlan_id	Unsigned integer	No			The IPv6 VLAN ID of a Grid member. Example: 10
nic_failover_enabled	Boolean	No		nic_failover_enabled	Set or deactivate the <b>nic_failover_enabled</b> flag of a Grid member. This is required only when the member is a single appliance. To enable this flag, you must set the <b>lan2_port</b> flag to <b>True</b> . Example: False
lan2_enabled	Boolean	No			Enable or disable LAN2. Example: False
lan2_ipv4addr	String	No		lan2_ipv4addr	The LAN IP address of the LAN2 port. To configure the LAN2 port, you must set the <b>lan2_port</b> flag to <b>True</b> . Example: 10.0.0.20
lan2_mask	String	No		lan2_mask	The netmask of the LAN2 port. To configure the LAN2 port, you must set the <b>lan2_port</b> flag to <b>True</b> . Example: 255.255.255.0
lan2_ipv6addr	String	No			The LAN IPv6 address of the LAN2 port. To configure the LAN2 port, you must set the <b>lan2_port</b> flag to <b>True</b> . Example: 2001::20
lan2_vlan_id	Unsigned integer	No			The VLAN ID of the IPv4 LAN2 port. To configure the LAN2 port, you must set the <b>lan2_port</b> flag to <b>True</b> . Example: 10
lan2_dscp	Unsigned integer	No			The DSCP value of the IPv4 LAN2 port. Valid values are integers between zero and 63. Example: 15
lan2_gateway	String	No		lan2_gateway	The gateway IP address of the LAN2 port. To configure the LAN2 port, you must set the <b>lan2_port</b> flag to <b>True</b> . Example: 10.0.0.1
lan2_ipv6_cidr	Unsigned integer	No		lan2_ipv6_cidr	The CIDR of the LAN2 port. To configure the LAN2 port, you must set the <b>lan2_port</b> flag to <b>True</b> . Example: 64

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
lan2_ipv6_dscp	Unsigned integer	No			The DSCP value of the IPv6 LAN2 port. Valid values are integers between zero and 63. Example: 15Example: 0
lan2_ipv6_enable	Boolean	No		lan2_ipv6_enable	Enable or disable IPv6 configuration of the LAN2 interface. Example: False
lan2_ipv6_enable_auto_config	Boolean	No		lan2_ipv6_enable_auto_config	Enable or disable IPv6 auto-configuration of the LAN2 interface.Example: False
lan2_ipv6_gateway	String	No		lan2_ipv6_gateway	The LAN IPv6 gateway of the LAN2 port. To configure the LAN2 port, you must set the <b>lan2_port</b> flag to <b>True</b> . Example: 2001::1
lan2_ipv6_vlan_id	Unsigned integer	No			The VLAN ID of the IPv6 LAN2 port. To configure the LAN2 port, you must set the <b>lan2_port</b> flag to <b>True</b> .Example: 20
nat_enabled	Boolean	No		nat_enabled	Specify <b>True</b> to enable the NAT compatibility setting or <b>False</b> to disable it. Example: False
nat_group	String	No		nat_group	The NAT group of a Grid member. This is required only when NAT compatibility is enabled. Example: group1/comm1,group2/comm2
nat_ip_address	String	No			The NAT IP address of a Grid member. Required only when NAT compatibility is enabled. Example: 10.0.0.10
static_routes	List of IPv4 static routes	No		static_routes	This list contains the static routes of a Grid member. The valid format is: address/subnet_mask/gateway. Example: 10.10.1.10/255.255.0.0/10.10.1.1
ipv6_static_routes	List of IPv6 static routes	No		ipv6_static_routes	This list contains the IPv6 static routes of a Grid member. The valid format is: address/cidr/gateway. Example: 2001::10/64/2001::1

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
remote_console_access	Boolean	No		use_remote_console_access	Enable or disable remote console access. When you enable remote console access, a client can access the device using a Secure Shell (SSH) connection. Example: False
support_access	Boolean	No		use_support_access	Specify <b>True</b> to set the <b>remote_console_access</b> flag or <b>False</b> to deactivate it. Default value is <b>True</b> . Example: False
enable_query	Boolean	No			Example: False
query_comm_string	String	No			Indicates the SNMP (Simple Network Management Protocol) community string that management systems must send together with their queries to the Infoblox appliance. If this is set to <b>undef</b> , SNMP queries will be disabled.
enable_snmpv3_query	Boolean	No		enable_snmpv3_query	Enable or disable SNMPv3 queries. Example: False
enable_traps	Boolean	No			Enable or disable traps. Example: False
trap_comm_string	String	No			The SNMP trap community string of the Grid member. Default value is <b>undef</b> .
enable_snmpv3_traps	Boolean	No			Enable or disable SNMPv3 traps. Example: False
snmp_admin	String	No			Set or retrieve the SNMP admin object. Example: sysname/ syscontact/syslocation/ sysdescr
snmpv3_query_users	String	No			Indicates the SNMPv3 user setting at the Grid level. Example: snmpv3User/comment
trap_receiver	String	No			Indicates the SNMP trap receivers of a Grid member. Example: 10.0.0.10/ snmpUserName/ comment

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
additional_ip_list	String	No		additional_ip_list	Indicates the additional IP list of a Grid member. This list contains additional interface information that can be used on the member level. The valid format is: interfaceType/networkSetting/comment/ ospf/anycast/bgp Example for IPv4: address/subnet_mask/gateway/vlan_id/primary/dscp/use_dscp Example for IPv6: virtual_ip/cidr_prefix/gateway/vlan_id/ primary/dscp/use_dscp/enabled/auto_router_configure_enabled
enable_member_redirect	Boolean	No		enable_member_redirect	Set or retrieve the flag that specifies if GUI redirection is enabled for members. Example: False
virtual_router_id	Integer	No			Indicates virtual router identifier. Example:110
dscp	Unassigned integer	No		dscp	Indicates the DSCP value. The default value is zero.
EA-Site	String	No	Extensible attribute	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. Example: California
EA-Users	String	No	Extensible attribute	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. Example: John
ADMGRP-XXXX	String	No	Permissions Admin Group/Role	permission	ADMGRP-JimSmith is an example of an admin permission of a specific admin group. Example: RW

## Upgrade Groups and Schedules

Consider the following when you import upgrade groups and their distribution and upgrade schedules:

- The appliance imports a new upgrade group only when both the distribution and upgrade schedules are inactive.
- You can modify members and the comment field of an existing upgrade group only when both the distribution and upgrade schedules are inactive.
- You can modify distribution related fields of an existing upgrade group only when the distribution schedule is inactive.
- You can modify upgrade related fields of an existing upgrade group only when the upgrade schedule is inactive.
- The Grid Master is the only member of the Grid Master group. You cannot move it to another upgrade group.
- You cannot change the members of the Reporting Member group nor move them to another upgrade group.
- When you remove a member from its original group, it will be placed in the Default group. You may notice additional members in the Default group if you remove members from any upgrade groups.

- When you specify both dependency and distribution or upgrade time for an upgrade group, dependency takes precedence.



**Note**

When you import an upgrade group and its distribution and upgrade schedules, you cannot control the activation and deactivation of the schedules. You can activate and deactivate the schedules through the GUI after the import.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-upgrade_group	String	Yes		UpgradeGroup	Example: UpgradeGroup
name	String	Yes	Name		Example: UpgradeGp
comment	String	No	Comment	comment	Add this field to overwrite the comment you entered for the upgrade group.
members	Grid member list	No	Member Assignment Name	members	Enter a list of Grid members separated by commas. Use the FQDNs of the Grid members. Enclose the entire string in double quotes. Example: "corp100.mktg.com,corp100.dev.com"
time_zone	String	No	Time Zone	time_zone	The value in this field applies to both distribution_time and upgrade_time. If you do not specify a time zone, the Grid level time zone is used.
distribution_dependent_group	String	No	Name	distribution_dependent_group	Enter the name of the preceding upgrade group that this group depends on during a distribution.
distribution_policy	String	No	Distribute to Members	distribution_policy	Valid values are: SIMULTANEOUSLY or SEQUENTIALLY.
distribution_time	Time	No	Start Distribution (Date/Time)	distribution_time	Enter the distribution start date and time in YYYY-MM-DDTHH:MM:SS format.
upgrade_dependent_group	String	No	Name	upgrade_dependent_group	Enter the name of the upgrade group that this group depends on during an upgrade.
upgrade_policy	String	No	Upgrade Members	upgrade_policy	Valid values are: SIMULTANEOUSLY or SEQUENTIALLY.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
upgrade_time	Time	No	Start Upgrade	upgrade_time	Enter the upgrade start date and time in YYYY-MM-DDTHH:MM:SS format.

## Named ACLs

Consider the following when you import named ACLs and ACEs for an existing named ACL:

- The appliance appends ACES to the end of the named ACL.
- Validate all imported named ACLs after a CSV import. The appliance does not automatically validate ACEs. To avoid conflicts and unexpected results, you must validate the named ACLs.
- To reorder an existing named ACL through CSV import, complete the following:
  - a. Export the named ACL.
  - b. Delete all ACEs in the named ACL. The appliance allows a named ACL without any ACEs.
  - c. Reorder the ACEs in the .csv file.
  - d. Re-import the updated file through CSV import.
- You can add a named ACL as a nested ACL to an existing named ACL.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-namedacl	String	Yes		NamedACL	Example: namedacl
name	String	Yes	Name	nacl	Example: nacl1
_new_name	String	No			Add this field to overwrite the name field when you select the overwrite or merge option.
comment	String	No	Comment	comment	Add this field to overwrite the comment you entered for the name ACL.

## Examples

This section contains examples of how to create data files for named ACLs. All examples use comma as the separator. You can use other supported separators, such as semicolon, space, or tab.

### Adding a new named ACL

This example shows how to import a new named ACL nacl1.

```
Header-namedacl,name*,comment namedacl,nacl1,"This is a new named ACL."
```

### Overwriting the name of a named ACL

This example shows how to overwrite the name of an existing named ACL nacl1 to nacl2.

```
Header-namedacl,name*,_new_name namedacl,nacl1,nacl2
```



## ACES in Named ACLs

You can add or overwrite ACEs in an existing named ACL. When you add or modify ACEs in a named ACL, you must have one of the following categories in each entry: IP address, TSIG key based ACE, or a nested named ACL. A combination of these in an entry will generate an error.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-namedaclitem	String	Yes		ACL item	Example: namedaclitem
parent	String	Yes	Named ACL	nacl	This is the existing named ACL. Example: nacl1
address	String	No	IPv4 Address IPv6 Address	ipv4addr ipv6addr	This field includes required information for IP address based ACEs. Use forward slashes as the delimiter to separate permission, address, and netmask. Example: "10.0.0.1/24/Allow"
_new_address	String	No			Add this field to overwrite the address field when you use the overwrite or merge option.
tsig_key	String	No	TSIG Key	TSIGKEY	This field includes required information or TSIG key based ACEs. Use forward slashes as the delimiter to separate tsig_key_name, tsig_key, tsig_key_alg, and use_2x_tsig_key. Example: "key_1/ny/ bY2Da8Lj+2YZ4dYEJLQ==/HMAC-SHA256/false"
_new_tsig_key	String	No			Add this field to overwrite the tsig_key field when you use the overwrite or merge option.
defined_acl	String	No	Named ACL	nacl	This field adds a named ACL as a nested ACL to an existing named ACL. Example: nacl2
_new_named_acl	String	No			Add this field to overwrite the named_acl field when you use the overwrite or merge option.
comment	String	No	Comment	comment	Add this field to overwrite the comment you entered for the name ACL.

### Adding an IP address based ACE to an existing named ACL

This example shows how to import an IP address based ACE to nacl1. You must include the following information in the address field: address/netmask/permission. The appliance append the ACE to the end of the ACL. Ensure that you validate the ACL after the import.

```
Header-namedaclitem,parent*,address namedaclitem,N1,10.0.38.230/ALLOW  
namedaclitem,N1,172.0.0.0/8/ALLOW
```

### Adding a TSIG key based ACE to an existing named ACL

This example shows how to import a TSIG key based ACE to nacl1. You must include the following information in the tsig\_key field: tsig\_key\_name/tsig\_key/tsig\_key\_alg/use\_2x\_tsig\_key. The appliance append the ACE to the end of the ACL. Ensure that you validate the ACL after the import.

```
Header-namedaclitem,parent*,tsig_key nmaedaclitem,"nacl1","key_1/  
bY2Da8Lj+2YZ4dYEJLQ==/HMAC-SHA256/false"
```

### Adding a nested named ACL to an existing named ACL

This example shows how to import a nested named ACL nacl2 to the parent named ACL nacl1. The appliance append the nested ACL to the end of the ACL. Ensure that you validate the ACL after the import.

```
Header-nmaedaclitem,parent*,defined_acl namedaclitem,"nacl1","nacl2"
```

### Adding and overwriting multiple ACEs

This example shows how to add new ACEs and modify existing ACEs. Ensure that you validate the ACL after the import.

```
Header-  
namedaclitem,parent*,address,_new_address,tsig_key,_new_tsig_key,defined_acl,  
_new_acl naemdaclitem,"nacl1","Allow/10.0.0.1/24","Deny/10.0.0.1/24",,,,  
namedaclitem,"nacl1",,,"key_1/bY2Da8Lj+2YZ4dYEJLQ==/HMAC-SHA256/false",,,  
namedaclitem,"nacl1",,,,,,"acl2"
```

## Discovery Credentials

You can define a list of SNMP v1/v2c credentials or a list of SNMPv3 credentials to import through a CSV file. Sensitive authentication information may be part of an SNMP credential import, including SNMPv1 or SNMPv2c communities or SNMPv3 user/password tuples for privacy and authentication. This information is also stored in the NIOS database. CSV import of credentials contains community strings and passwords in plain text. Imported CSV files are uploaded with a POST in an HTTPS session and are deleted immediately after the import operation completes, whether or not the import is successful.

### SNMPv1/SNMPv2c Credentials Format

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guideline
Header-discovery SNMPv1v2 credentials	String	Yes		SNMPCredential	
comment	String	No	Comment	comment	Use this field to set or retrieve the comment on SNMPv1 and SNMPv2 users.
community_string	String	Yes	Read Community	community_string	
_new_community_string	String	No			Use this field to set or retrieve the public community string.
parent	String	Yes	SNMPv1/v2		This is the existing credential.

### SNMPv3 Credentials Format

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guideline
Header-discovery SNMPv3 credentials	String	Yes		SNMP3Credential	
user	String	Yes	Name	user	Use this field to set or retrieve the SNMPv3 user name.
_new_user	String	No	Name		Add this field to overwrite the user field when you use the overwrite or merge option.
authentication_password	String (if protocol is (MD5/SHA))	No	Auth Password	authentication_password	Use this field to set the authentication password if the authentication protocol is MD5 or SHA. This is a write-only attribute.
authentication_protocol	String (valid values are MD5/SHA/NONE/SHA-224/SHA-256/ SHA-384/SHA-512)	No	Auth Protocol	authentication_protocol	Use this field to set or retrieve the authentication protocol.
_new_authentication_protocol	String	No			Add this field to overwrite the authentication_protocol field when you use the overwrite or merge option.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guideline
comment	String	No	Comment		Add this field to overwrite the comment you entered for the credential.
parent	String	Yes	SNMPv3		
privacy_password	String (if protocol is AES/DES/3DES)	No	Privacy Password	privacy_password	Use this field to set the privacy password if the privacy protocol is 'AES', 'DES' or '3DES'. This is a write-only attribute.
privacy_protocol	String (valid values: AES-128/DES/3DES/NONE/AES-192/AES-192C/AES-256/AES-256C)	No	Privacy Protocol	privacy_protocol	Use this field to set or retrieve the privacy protocol.
_new_privacy_protocol	String	No			Add this field to overwrite the privacy_protocol field when you use the overwrite or merge option.

## DTC Header Items

This topic lists all the DTC objects whose names start with "dct". These objects are displayed as DtcLbdn, DtcServer, and so on in the CSV export or import file. The object DtcCertificate supports only CSV export, whereas other objects support CSV import and export. For details, see the following sections:

- [DtcLbdn Header Items](#)
- [DtcServer Header Items](#)
- [DtcPool Header Items](#)
- [DtcMonitorHttp Header Items](#)
- [DtcMonitorSip Header Items](#)
- [DtcMonitorIcmp Header Items](#)
- [DtcMonitorPdp Header Items](#)
- [DtcMonitorTcp Header Items](#)
- [DtcMonitorSnmp Header Items](#)
- [DtcARecord Header Items](#)
- [DtcAaaaRecord Header Items](#)
- [DtcCnameRecord Header Items](#)
- [DtcNaptrRecord Header Items](#)
- [DtcCertificate Header Items](#)

### DtcLbdn Header Items

Name	Type	Required	Example	ibap Name	Comment
Header-DtcLbdn	String	Yes	DtcLbdn		

Name	Type	Required	Example	lbap Name	Comment
name*	String	Yes	load_bal		
_new_name	String	No		name	
lb_method*	String	Yes	GLOBAL_AVAILABILITY	lb_method	Valid values are 'GLOBAL_AVAILABILITY', 'RATIO', 'ROUND_ROBIN', 'TOPOLOGY' and 'SOURCE_IP_HASH'
patterns	List of strings	No	.domain.com, .test.com'	patterns	Valid value is an array of FQDN patterns in string format
disabled	Boolean	No	false	disabled	
comment	String	No	A DTC LBDN comment	comment	
persistence	Unsigned integer	No	5	persistence	Zero specifies no caching
topology	String	No	topology-1	topology	
ttl	Unsigned integer	No	10	ttl	
pools	LBDN linked list of pools	No	pool/ratio eg: "pool1/1,pool2/10...	pools	
auth_zones	LBDN linked list of auth zones	No	Must be in the zone_name/ view_name format. For example:  authzone1.com/view1, authzone2.com/view2...	auth_zones	List of DTC LBDN linked authoritative zones
types	String	No	NAPTR	types	Valid values are 'A', 'AAA', 'NAPTR', 'CNAME'
priority	Unsigned integer	No	3	priority	Valid values between 1 and 3. The default is 3.
EA-Site	String	No	San Jose		

### DtcServer Header Items

Name	Type	Required	Example	lbap Name	Comment
Header-DtcServer	String	Yes	DtcServer		
name*	String	Yes	server1	name	
_new_name	String	No		name	

Name	Type	Required	Example	Ibap Name	Comment
host*	String	Yes	192.168.1.2	host	Address or FQDN of a DTC server
disabled	Boolean	No	false	disabled	
comment	String	No	A DTC server comment	comment	
auto_create_host_record	Boolean	No	false	comment	
sni_hostname	String	No	sni-host.infoblox.localdomain	sni_hostname	Host name for Server Network Indication to be used with the HTTPS monitor
monitors	List of DTC health monitors	No	Monitor name/Monitor type/Host For example: TCP-1/tcp/10.10.10.10,SIP-1/sip/20.20.20.20,ICMP-1/icmp/30.30.30.30,...	monitors	List of names of DTC server monitors
EA-Site	String	No	San Jose		

#### DtcPool Header Items

Name	Type	Required	Example	Ibap Name	Comment
Header-DtcPool	String	Yes	DtcPool		
name*	String	Yes	pool1	name	
_new_name	String	No		name	
lb_preferred_method	String	No	RATIO	lb_preferred_method	Valid values are 'ALL_AVAILABLE', 'DYNAMIC_RATIO', 'GLOBAL_AVAILABILITY', 'RATIO', 'ROUND_ROBIN', 'TOPOLOGY', and 'SOURCE_IP_HASH'
disabled	Boolean	No	false	disabled	
comment	String	No	A DTC pool comment	comment	
lb_preferred_topology	Topology rules for preferred 'TOPOLOGY' load balancing method	No	custom-2	preferred_topology	

Name	Type	Required	Example	lbap Name	Comment
lb_alternate_method	String	No	RATIO	alternate_method	Alternate load balancing method. Valid values are 'ALL_AVAILABLE', 'DYNAMIC_RATIO', 'NONE', 'GLOBAL_AVAILABILITY', 'RATIO', 'ROUND_ROBIN', 'TOPOLOGY', and 'SOURCE_IP_HASH'.
lb_alternate_topology	Topology rules for alternate 'TOPOLOGY' load balancing method	No	custom-3	alternate_topology	
availability	String	No	any	availability	DTC pool resources availability status. Valid values are 'ALL', 'ANY', and 'QUORUM'.
quorum	Unsigned integer	No	20	quorum	The number of monitors that must report the resource as 'up' for 'QUORUM' availability mode to be available.
ttl	Unsigned integer	No	10	ttl	
servers	List of DTC servers	No	server/ratio For example: dtc-server1/1,dtc-server2/10...	servers	
monitors	List of health monitors	No	Monitor name/Monitor type For example: ICMP-1/icmp,HTTP-1/http,...	monitors	List of names and monitor types of DTC pool monitors
dynamic_ratio_preferred	Preferred dynamic ratio load balancing settings	No	method/monitor/monitor_metric/ monitor_weighting/ invert_monitor_metric For example: MONITOR/snmp.1.3/ RATIO/false	dynamic_ratio_preferred	When preferred_method="DYNAMIC_RATIO"
dynamic_ratio_alternate	Alternate dynamic ratio load balancing settings	No	method/monitor/monitor_metric/ monitor_weighting/ invert_monitor_metric For example: MONITOR/snmp.1.3/ PRIORITY/false	dynamic_ratio_alternate	When preferred_method="TOPOLOGY" and alternate_method="DYNAMIC_RATIO"
EA-Site	String	No	San Jose		

## DtcMonitorHttp Header Items

Name	Type	Required	Example	lbap Name	Comment
Header-DtcMonitorHttp	String	Yes	DtcMonitorHttp		
name*	String	Yes	http_monitor1	name	
_new_name	String	No		name	
comment	String	No	A DTC HTTP monitor comment	comment	
interval	Unsigned integer	No	10	interval	
timeout	Unsigned integer	No	7	timeout	Valid values between 1 and 15
port	Unsigned integer	No	8080	port	Valid values between 1 and 65535
request	String	No	GET /	request	Maximum of 1024 characters
content_check	String	No	EXTRACT	content_check	Valid values are 'NONE'(default), 'MATCH', and 'EXTRACT'. If 'MATCH', then 'content_check_op' and 'content_check_regex' params are required; if 'EXTRACT' then 'content_check_op', 'content_extract_type', 'content_extract_value', and 'content_check_regex' values are required.
content_check_input	String	No	BODY	content_check_input	Valid values are 'HEADERS', 'ALL'(default) and 'BODY'
content_check_regex	String	No	SQL Error	content_check_regex	
content_check_op	String	No	EQ	content_check_op	Valid values are 'EQ' and 'NEQ' for 'MATCH' content check, and 'EQ', 'NEQ', 'LEQ', and 'GEQ' for 'EXTRACT' content check type
content_extract_group	Unsigned integer	No	3	content_extract_group	Valid values between 0 and 8. The default is 0.
content_extract_type	String	No	INTEGER	content_extract_type	Valid values are 'INTEGER' and 'STRING'. The default is 'STRING'.
content_extract_value	String	No	1	content_extract_value	A desired extraction value in string format



Name	Type	Required	Example	lbap Name	Comment
result	String	No	CODE_IS	result	Valid values are 'ANY', 'CODE_IS', and 'CODE_IS_NOT'
result_code	Unsigned integer	No	300	result_code	Valid values between 0 and 999
secure	Boolean	No	true	secure	Default is 'false'
client_cert	Client certificate	No	626596e ... 4362f80c (128 characters)	client_cert	Valid value is a DTC certificate object.
ciphers	String	No	DHE-RSA-AES256-SHA	ciphers	Valid value is array of ciphers in a string format
retry_up	Unsigned integer	No	3	retry_up	Valid values between 1 and 10
retry_down	Unsigned integer	No	3	retry_down	Valid values between 1 and 10
validate_cert	Boolean	No	false	validate_cert	Default is 'true'
enable_sni	Boolean	No	true	enable_sni	Default is 'false'
EA-Site	String	No	San Jose		

### DtcMonitorSip Header Items

Name	Type	Required	Example	lbap Name	Comment
Header-DtcMonitorSip	String	Yes	DtcMonitorSip		
name*	String	Yes	sip_monitor1	name	
_new_name	String	No		name	
comment	String	No	A DTC SIP monitor comment	comment	
interval	Unsigned integer	No	10	interval	
timeout	Unsigned integer	No	7	timeout	Valid values between 1 and 15

Name	Type	Required	Example	lbap Name	Comment
port	Unsigned integer	No	8080	port	Valid values between 1 and 65535
request	String	No	GET /	request	Maximum of 1024 characters
result	String	No	CODE_IS	result	Valid values are 'ANY', 'CODE_IS', and 'CODE_IS_NOT'
result_code	Unsigned integer	No	300	result_code	Valid values between 0 and 999
transport	String	No	UDP	transport	Valid values are 'TCP', 'UDP', 'SIPS', and 'TLS'
client_cert	Client certificate	No	626596e ... 4362f80c (128 characters)	client_certificate	Valid value is a DTC certificate object
ciphers	String	No	DHE-RSA-AES256-SHA	ciphers	Valid value is array of ciphers in a string format
retry_up	Unsigned integer	No	3	retry_up	Valid values between 1 and 10
retry_down	Unsigned integer	No	3	retry_down	Valid values between 1 and 10
validate_cert	Boolean	No	false	validate_certificate	Default is 'true'
EA-Site	String	No	San Jose		

### DtcMonitorIcmp Header Items

Name	Type	Required	Example	lbap Name	Comment
Header-DtcMonitorIcmp	String	Yes	DtcMonitorIcmp		
name*	String	Yes	icmp_monitor1	name	
_new_name	String	No		name	
comment	String	No	A DTC ICMP monitor comment	comment	
interval	Unsigned integer	No	10	interval	
timeout	Unsigned integer	No	7	timeout	Valid values between 1 and 15
retry_up	Unsigned integer	No	3	retry_up	Valid values between 1 and 10

Name	Type	Required	Example	lbap Name	Comment
retry_down	Unsigned integer	No	3	retry_down	Valid values between 1 and 10
EA-Site	String	No	San Jose		

#### DtcMonitorPdp Header Items

Name	Type	Required	Example	lbap Name	Comment
Header-DtcMonitorPdp	String	Yes	DtcMonitorPdp		
name*	String	Yes	pdp_monitor1	name	
_new_name	String	No		name	
comment	String	No	A DTC PDP monitor comment	comment	
interval	Unsigned integer	No	10	interval	
timeout	Unsigned integer	No	7	timeout	Valid values between 1 and 15
retry_up	Unsigned integer	No	5	retry_up	Valid values between 1 and 10
retry_down	Unsigned integer	No	5	retry_down	Valid values between 1 and 10
port	Unsigned integer	No	6030	port	Valid values between 1 and 65535
EA-Site	String	No	San Jose		

#### DtcMonitorTcp Header Items

Name	Type	Required	Example	lbap Name	Comment
Header-DtcMonitorTcp	String	Yes	DtcMonitorTcp		
name*	String	Yes	tcp_monitor1	name	
port*	Unsigned integer		6030	port	Valid values between 1 and 65535

Name	Type	Required	Example	lbap Name	Comment
_new_name	String	No		name	
comment	String	No	A DTC PDP monitor comment	comment	
interval	Unsigned integer	No	10	interval	
timeout	Unsigned integer	No	7	timeout	Valid values between 1 and 15
retry_up	Unsigned integer	No	3	retry_up	Valid values between 1 and 10
retry_down	Unsigned integer	No	3	retry_down	Valid values between 1 and 10
EA-Site	String	No	San Jose		

#### DtcMonitorSnmp Header Items

Name	Type	Required	Example	lbap Name	Comment
Header-DtcMonitorSnmp	String	Yes	DtcMonitorSnmp		
name*	String	Yes	snmp_monitor1	name	
port	Unsigned integer		6030	port	Valid values between 1 and 65535
_new_name	String	No		name	
comment	String	No	A DTC SNMP monitor comment	comment	
interval	Unsigned integer	No	10	interval	
timeout	Unsigned integer	No	7	timeout	Valid values between 1 and 15
retry_up	Unsigned integer	No	3	retry_up	Valid values between 1 and 10
retry_down	Unsigned integer	No	3	retry_down	Valid values between 1 and 10

Name	Type	Required	Example	Ibap Name	Comment
version	String	No	V1	version	Valid values are 'V1', 'V2C', and 'V3'
community	String	No	desired_community	commu nity	SNMP community string for an SNMP authentication
oids	List of OIDs for SNMP monitoring	No	.1.3/Comment-1/INTEGER/RANGE/10/1000,.1.2/Comment-2/STRING/EXACT/abc/...	oids	
user	String	No	user1	user	SNMPv3 user setting
context	String		desired_context	context	SNMPv3 context. Maximum of 1023 characters.
engine_id	String	No	desired_engine_id	engine_ id	SNMPv3 engine identifier. Max of 1023 characters.
EA-Site	String	No	San Jose		

#### DtcARecord Header Items

Name	Type	Required	Example	Ibap Name	Comment
Header-DtcARecord	String	Yes	DtcARecord		
ipv4addr*	IPv4 address	Yes	10.0.0.1	address	
_new_ipv4addr	IPv4 address	No		address	
dtc_server*	DTC server	Yes	server1	lbdns_server	DTC server the DTC A record is associated with
ttl	Unsigned integer	No	1024	ttl	
disabled	Boolean	No	false	disabled	
comment	String	No	Sample DTC A record	comment	

#### DtcAaaaRecord Header Items

Name	Type	Required	Example	Ibap Name	Comment
Header-DtcAaaaRecord	String	Yes	DtcAaaaRecord		

Name	Type	Required	Example	lbap Name	Comment
ipv6addr*	IPv6 address	Yes	2001:db8::1	address	
_new_ipv6addr	IPv6 address	No		address	
dtc_server	DTC server	Yes	server1	dtc_server	DTC Server the DTC AAAA record is associated with
ttd	Unsigned integer	No	1024	ttd	
disabled	Boolean	No	false	disabled	
comment	String	No	Sample DTC AAAA record	comment	

### DtcCnameRecord Header Items

Name	Type	Required	Example	lbap Name	Comment
Header-DtcCnameRecord	String	Yes	DtcCnameRecord		
canonical*	String	Yes	dtc.localdomain	canonical_name	Canonical name of the resource
_new_canonical	String	No		canonical_name	
dtc_server*	DTC server	Yes	server1	dtc_server	DTC Server the DTC CNAME record is associated with
ttd	Unsigned integer	No	1024	ttd	
disabled	Boolean	No	false	disabled	
comment	String	No	Sample DTC CNAME record	comment	

### DtcNaptrRecord Header Items

Name	Type	Required	Example	lbap Name	Comment
Header-DtcNaptrRecord	String	Yes	DtcNaptrRecord		
dtc_server*	DTC server	Yes	server1	dtc_server	DTC server the DTC NAPTR record is associated with

Name	Type	Required	Example	lbap Name	Comment
order*	Unsigned integer	Yes	100	order	
_new_order	Unsigned integer	No		order	
preference*	Unsigned integer	Yes	10	preference	
_new_preference	Unsigned integer	No		preference	
replacement*	String	Yes	domain2.com	replacement	The desired replacement value in a Fully-Qualified Domain Name (FQDN) format.
_new_replacement	String	No		replacement	
flags*	String	No	U	flags	Empty value is allowed. Supported values for the flags field are 'U', 'S', 'P', and 'A'.
_new_flags	String	No		flags	
services*	String	No	SIP+D2U	services	Empty value is allowed
_new_services	String	No		services	
regex*	String	No	!http://my[.](.*)!1i	regex	Empty value is allowed
_new_regex	String	No		services	
ttl	Unsigned integer	No	28800	ttl	
disabled	Boolean	No	false	disabled	
comment	String	No	Sample DTC NAPTR record	comment	

#### DtcCertificate Header Items

Name	Type	Required	Example	lbap Name	Comment
Header-Dtcertificate	String	Yes	DtcCertificate		

Name	Type	Required	Example	lbap Name	Comment
certificate	DTC certificate	No	Object Hash/Issuer/Valid from/Valid to/Subject For example: "702971b9b6bb34468f6006389f53849dff43c78ee415d256f771cdcb58782081025e28362c714d27c9c652967afa64f0edf0c17170b3ace72bd0b3c9eebad7ef/CN=""pannpn"", OU=""Engineering"", O=""Infoblox"", L=""NYC", ST=""Kerala"" C=""IN""/2017-11-22 18:46:43/2018-11-22 18:46:43/CN=""pannpn"", OU=""Engineering"", O=""Infoblox"", L=""NYC"", ST=""Kerala"", C=""IN""	cert	Object hash is the SHA512 fingerprint of the certificate, that is to be used during the CSV import/export of dtcMonitorHttp and dtcMonitorSip objects in the client_cert field

## DHCP Lease



### Note

Infoblox supports CSV export for DHCP lease, but does not support CSV import.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
header-lease	String	Yes			Example: Lease
agent_id	String				Determined based on the value of "option"
billing_class	String			billing_class	
binding_state	Enum{FREE, BACKUP, ACTIVE, EXPIRED, RELEASED, ABANDONED, RESET}			binding_state	
circuit_id	String	No		agent_circuit_id	This field is required when match_option = CIRCUIT_ID. Example: 11
duid	String	Yes	DUID	duid	Example: 0001
ends	DateTime			ends	Lease end time
fingerprint	String	Yes	DHCP Fingerprints	fingerprint	An array of DHCP fingerprints.
fingerprint-class	String	No	Next Server	next_server	Example: blue.domain.com



Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ip_address	IP address	Yes	IP Address	ipv4addr	Example: 10.0.0.11
link_selection	IP address		IP address		Example: 10.0.0.0
mac_address	MAC address	No	MAC Address	mac	This field is required if match_option = MAC_ADDRESS, or if you are importing a reservation. Example: aa:d1:dd:10:ff:02
next_binding_state	Enum{FREE, BACKUP, ACTIVE, EXPIRED, RELEASED, ABANDONED, RESET}			next_binding_state	Example: FREE
protocol	String	Yes	protocol	protocol	Protocol type. This can be IPV4 or IPV6.
remote_id	Integer	No	Remote ID	remote_id_name	Example: 50
server_host_name	String	No			Indicates the name of the server host. Example: infoblox.localdomain
server_id_override	IP address		IP address		Example: 10.0.0.1
starts	DateTime				Lease start time
uid	String	Yes	UID	uid	
username	String		Username	username	

## Subscriber Record

Field Name	Data Type	Required (Yes/No)	Usage and Guidelines	Example
ip_addr	IP address	Yes	Can be IPv4 or IPv6	2002::abcd or 20.20.20.20
prefix	Integer	Yes	<ul style="list-style-type: none"> <li>For IPv6, prefix must not be greater than 128</li> <li>For IPv4, prefix must not be greater than 32</li> </ul>	<ul style="list-style-type: none"> <li>32 if ip_addr is 20.20.20.20</li> <li>128 if ip_addr is 2002::abcd</li> </ul>

Field Name	Data Type	Required (Yes/No)	Usage and Guidelines	Example
localid	String	Yes	<ul style="list-style-type: none"> <li>Must be a valid hexadecimal string if it is MAC_Addr</li> <li>Can be N/A if not available</li> </ul>	a1b1c1d0e0f0 or N/A
ipsd	String	Yes	<ul style="list-style-type: none"> <li>Can be N/A if not available</li> <li>Can be a port number in case of CGNAT support</li> </ul>	BLR or NYK
flags	String	Yes	<ul style="list-style-type: none"> <li>S for static record</li> <li>B for black and white list support</li> <li>SB for both static and blacklist and whitelist support</li> </ul>	S or B or SB
alt_ip_addr	IP address	No	<ul style="list-style-type: none"> <li>Must belong to a different type than ip_addr</li> <li>Prefix can be a part of alt_ip_addr. Only 32 and 128 are supported as prefix values</li> </ul>	<ul style="list-style-type: none"> <li>2002::abcd if ip_addr is 20.20.20.20</li> <li>20.20.20.20 if ip_addr is 2002::abcd</li> <li>2002::abcd/128</li> <li>20.20.20.20/32</li> </ul>
subscriber_id	String	No	Must be in the avp_name = avp_value format	<ul style="list-style-type: none"> <li>User_Name = Steve</li> <li>IMEI = 123456</li> <li>IMSI = 78910</li> </ul>
site	String	No	Site same must be the same as that configured on the <b>Subscribers Site</b> tab in Grid Manager	pc_site1
nas_contextual	String	No	<ul style="list-style-type: none"> <li>Must be in the avp_name = avp_value format</li> <li>avp_name must be one among the list of Attribute Value Pairs displayed on the <b>NAS Contextual Information</b> list on the <b>Subscriber Services Properties</b> tab</li> </ul>	<ul style="list-style-type: none"> <li>NAS-PORT = 25000</li> <li>NAS-IP-Address = 20.20.20.20</li> </ul>

Field Name	Data Type	Required (Yes/No)	Usage and Guidelines	Example
ancillaries	String	No	<ul style="list-style-type: none"> <li>Must be in the avp_name = avp_value format</li> <li>avp_name must be one among the list of Attribute Value Pairs displayed on the <b>Ancillary Fields</b> list on the <b>Subscriber Services Properties</b> tab</li> </ul>	Subscriber-Secure-Policy=F0FB or Proxy-All=1
accounting_session_id	String	No	<ul style="list-style-type: none"> <li>Must be an alphanumeric string</li> </ul>	Client-408123-8817 or Acct-Session-Id=408123-8817
subscriber_secure_policy	String	No	<ul style="list-style-type: none"> <li>Must be a valid hexadecimal string</li> <li>Length must not be greater than 8 because it is a 32-bit binary value</li> <li>Length of the hexadecimal string must be an even number</li> </ul>	0x123456 or 0xABCD or 0x12EFB7 or FFEE4321 or ABCDEF12
parental_control_policy	String	No	<ul style="list-style-type: none"> <li>Must be a valid hexadecimal string</li> <li>Length must not be greater than 32 because it is a 128-bit binary value</li> <li>Length of the hexadecimal string must be an even number</li> </ul>	123FAB78FA987654F000103A0B201245
wpc_category_policy	String	No	<ul style="list-style-type: none"> <li>Must be a valid hexadecimal string</li> <li>Length must not be greater than 32 because it is a 128-bit binary value</li> <li>Length of the hexadecimal string must be an even number</li> </ul>	123FAB78FA987654F000103A0B201245
unknown_category_policy	Boolean	No	Must be either TRUE or FALSE	TRUE or FALSE
dynamic_category_policy	Boolean	No	Must be either TRUE or FALSE	TRUE or FALSE
bwflag	Boolean	No	Must be either TRUE or FALSE	TRUE or FALSE
proxy_all	Boolean	No	Must be either TRUE or FALSE	TRUE or FALSE

Field Name	Data Type	Required (Yes/No)	Usage and Guidelines	Example
black_list	String	No	<ul style="list-style-type: none"> <li>bwflag must be enabled to specify a value for this field</li> <li>Must have a list of valid domain names separated by a comma</li> </ul>	facebook.com or Facebook.com,bad.com,verybad.com
white_list	String	No	<ul style="list-style-type: none"> <li>bwflag must be enabled to specify a value for this field</li> <li>Must have a list of valid domain names separated by a comma</li> </ul>	google.com or google.com,good.com,verygood.com

## Member DHCP Objects

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-MemberDhcp	String	Yes			Identifies the first row as a header row for the member DHCP objects. Example: MemberDhcp
broadcast_address	IP address	No	Broadcast Address		This field applies to the host address. Example: 10.1.0.
domain_name_servers	IP address list	No	Name server		List of domain name servers. Example: "10.2.3.4,11.2.3.4,"
ignore_client_requested_options	Boolean	No			Clears the value of option-55 when you set the value to True. Example: True
pxe_lease_time	Unsigned integer	No	PXE Lease Time	pxe_lease_time	Indicates the lease time for PXE clients in seconds. This field applies to the host address. Example: 43220.
lease_time	Unsigned integer	No	Lease Time	lease_time	Indicates the <b>lease_time</b> attribute of a DHCP NAC filter object. This field applies to the host address.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
domain_name	FQDN	No	Domain Name	domain_name	Indicates the domain name. This field applies to the host address.
routers	IP address list	No	Router		Indicates the router IP addresses. This field applies to the host address. Example: "2.0.0.2,1.2.3.4"
option_logic_filters	List of IPv4 logic filter rules		Filter Type/Action		Examples: .com.infoblox.dns.dhcp_mac_filter\$mac_filter_name, .com.infoblox.dns.nac_filter\$nac_filter_name, .com.infoblox.dns.dhcp_option_filters\$opt_filter_name
enable_pxe_lease_time	Boolean	No	Enable PXE Lease Time		If this value is set to <b>True</b> , DHCP server uses different lease time for PXE clients. This field applies to the host address. Example: False
deny_bootp	Boolean	No	Deny BOOTP Requests	deny_bootp	When this is set to <b>True</b> , it denies BOOTP requests. This field applies to the host address. Example: FALSE
bootfile	String	No	Boot File	bootfile	Indicates the boot file name. Example: bootfile1
bootserver	String	No	Boot Server	bootserver	Indicates the boot server. Example: abc.corp100.com
nextserver	String	No	Next Server	nextserver	Indicates the next server. Example: blue.domain.com

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
enable_thresholds	Boolean	No	Enable DHCP Thresholds	enable_dhcp_thresholds	Enable DHCP thresholds. When you set this field to <b>True</b> , you must enter values in the <b>range_high_water_mark</b> and <b>range_low_water_mark</b> fields. You cannot leave these fields empty. Otherwise, the appliance displays an error message.
range_high_water_mark	Unsigned integer	No	High Water Mark	high_water_mark	Indicates the percentage value for DHCP range usage after which an alarm will be active. When you set <b>enable_thresholds</b> to <b>True</b> , you must enter values in this field and in the <b>range_low_water_mark</b> field. You cannot leave these fields empty. Otherwise, the appliance displays an error message. Example: 80
range_high_water_mark_reset	Unsigned integer	No			Indicates the percentage value for DHCP range usage after which an alarm will be reset. Example: 85
range_low_water_mark	Integer	No	Low Water Mark	low_water_mark	Indicates the percentage value for DHCP range usage below which an alarm will be active. When you set <b>enable_thresholds</b> to <b>True</b> , you must enter values in this field and in the <b>range_high_water_mark</b> field. You cannot leave these fields empty. Otherwise, the appliance displays an error message. Example: 10

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
range_low_water_mark_reset	Unsigned integer	No			Indicates the percentage value for DHCP range usage above which an alarm will be reset. Example: 10
enable_threshold_email_warnings	Boolean	No	Enable Email Warnings	enable_email_warnings	When you use <b>Merge</b> , the appliance preserves the existing value in this field. When you use <b>Override</b> , you must include a value, either <b>True</b> or <b>False</b> . Otherwise, the appliance displays an error message. Example: <b>True</b>
enable_threshold_snmp_warnings	Boolean	No	Enable SNMP Warnings	enable_snmp_warnings	Send DHCP threshold warnings via SNMP. When you use <b>Merge</b> , the appliance preserves the existing value in this field. When you use <b>Override</b> , you must include a value, either <b>True</b> or <b>False</b> . Otherwise, the appliance displays an error message. Example: TRUE
threshold_email_addresses	Email address list	No	Email Addresses	email_list	List of email addresses. Example: "admin1@infoblox.com", "admin2@somewhere.com"
enable_ddns	Boolean	No	Enable DDNS Updates	enable_ddns	Enable or disable dynamic updates via DHCP to DNS server(s). Example: FALSE
enable_option81	Boolean	No	Option 81 Support	ddns_use_option81	Enable or disable option 81 support. Enables <b>always_update_dns</b> field. Example: TRUE
always_update_dns	Boolean	No	FQDN Support	ddns_server_always_updates	Updates DNS when the value is <b>True</b> . Example: TRUE

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
generate_hostname	Boolean	No	Generate Hostname	override_ddns_generate_hostname	Generates host name only if not sent by client when the value is <b>True</b> . Does not generate host name when the value is <b>False</b> .
update_static_leases	Boolean	No	Fixed Address Updates	ddns_update_fixed_address	When the value is set to <b>True</b> , DHCP server will update DNS for client with static IP address.
ddns_ttl	Unsigned integer	No	DDNS Update TTL	ddns_ttl	Indicates the DDNS TTL value in seconds. This is an inherited field. Example: 1200
update_dns_on_lease_renewal	Boolean	No	Lease Renewal Update	override_update_dns_on_lease_renewal	Indicates whether the DHCP server updates DNS when a DHCP lease is renewed. Specify <b>True</b> to enable it or <b>False</b> to disable it.
preferred_lifetime	Unsigned integer	No	Preferred Lifetime	preferred_lifetime	Indicates whether the <b>preferred_lifetime</b> value in the DHCP member is used, instead of the Grid default. Example: 604800
valid_lifetime	Unsigned integer	No	Valid Lifetime	valid_lifetime	Indicates whether the <b>valid_lifetime</b> value in the Grid Member DHCP is used, instead of the upper-level default. Example: 43200
name	Grid Member	Yes			The name of the Grid member. Example: member.infoblox.com
is_authoritative	Boolean	No	Authoritative	authority	Set this value to <b>True</b> to perform override operation. Example: FALSE



Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
recycle_leases	Boolean	No	Lease Deletion	recycle_leases	This field is set to <b>True</b> by default. When you set this to <b>True</b> , leases in a deleted range are kept until expiration. Ensure that you use the <b>Override</b> option if you want to change the value to <b>False</b> . Merging data from an import preserves the default value.
ping_count	Unsigned integer	No		ping_count	Indicates the number of DHCP pings. Example: 1
ping_timeout	Unsigned integer	No		ping_timeout	Indicates the timeout (in seconds) for DHCP pings. Example: 1000
enable_leasequery	Boolean	No			When you set this value to <b>True</b> , the appliance allows lease query. Example: False
retry_ddns_updates	Boolean	No		retry_ddns_updates	When you set this value to <b>True</b> , DHCP server will retry failed DNS updates. Example: False
ddns_retry_interval	Unsigned integer	No		ddns_retry_interval	Indicates the minimum time in minutes between DNS update retries. You must set <b>ddns_retry_updates</b> to <b>True</b> to modify the <b>ddns_retry_interval</b> value in the CSV file.
lease_scavenge_time	Unsigned integer	No			Indicates the Grid level <b>lease_scavenge_time</b> value. If the value is -1, which means this lease scavenge will be disabled. The minimum value would be $7 * 24 * 60 * 60$ ( 7 days).

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
enable_fingerprint	Boolean	No			When you set this value to <b>True</b> , fingerprint matching for incoming lease requests will be enabled. Example: False.
ipv6_enable_ddns	Boolean	No			Set this value to <b>True</b> to override the value at the Grid level. Set the parameter to <b>False</b> to inherit the settings from the Grid. Example: False
ipv6_ddns_enable_option_fqdn	Boolean	No			Indicates whether the FQDN option sent by the client is to be used, or if the server can automatically generate the FQDN. Default value is half of lease time. Example: False
ipv6_generate_hostname	Boolean	No			When you set this value to <b>True</b> , the hostname is generated if it is not sent by the client. Example: False
ipv6_ddns_domainname	String	No			The DDNS domain name in FQDN format. Example: test_domain.com
ipv6_ddns_ttl	Unsigned integer	No			Indicates the member IPv6 DDNS TTL value in seconds. Example: 0
ipv6_domain_name_servers	IPv6 DNS server list	No			List of IPv6 domain name servers. Example: "2001::1, 2001::2..."
ipv6_domain_name	String	No			To override the upper-level setting for <b>ipv6_domain_name</b> , you must set the value to <b>True</b> . Set the parameter to <b>False</b> to inherit the upper-level setting for ipv6_domain_name.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ipv6_recycle_leases	Boolean	No			When you set this to <b>True</b> , the leases are kept in recycle bin until one week after expiration. When you set this to <b>False</b> , the leases are irrecoverably deleted. Example: False
ipv6_server_duid	String	No			DUID (DHCPv6 unique identifier) of the Grid member in string format.
ipv6_enable_retry_updates	Boolean	No			When you set this flag to <b>True</b> , the DHCPv6 server retries failed dynamic DNS updates. The default value is <b>True</b> . Example: False
ipv6_retry_updates_interval	Unsigned integer	No			Set the retry interval when the member DHCPv6 server makes repeated attempts to send DDNS updates to a DNS server. The default retry interval is five minutes.
ipv6_update_dns_on_lease_renewal	Boolean	No			Set or retrieve the <b>ipv6_update_dns_on_lease_renewal</b> flag. This attribute controls whether the DHCP server updates DNS when an IPv6 DHCP lease is renewed. Specify <b>True</b> to enable this feature or <b>False</b> to disable it. The default value is <b>False</b> .
ddns_domainname	String	No			The DDNS domain name in FQDN format. Example: test_domain.com

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
lease_per_client_settings	String	No			Defines how the server will release the client lease. This field is set to <b>RELEASE_MATCHING_ID</b> by default. Valid values are ONE_LEASE_PER_CLIENT, RELEASE_MATCHING_ID and NEVER_RELEASE.
ignore_client_identifier	Boolean	No			Indicates if the client identifier will be ignored for a DHCP shared network object. When you set this to <b>True</b> , the client identifier will be ignored. Example: False
OPTION-1	String	No	Custom DHCP Options	options	This field applies to the host address. Example: '255.0.0.0' name implies vendor_class='DHCP' (default)
OPTION-XXXX-200	Option information	No	Custom DHCP Options	options	This field applies to the host address. Example: 'dfdfdf' name implies vendor_class='XXXX', optioncode/number 200
ADMGRP-XXXX	String	No	Permissions Admin Group/Role	permission	ADMGRP-JimSmith is an example of an admin permission of a specific admin group. Example: RW
v6_leases_scavenging_enabled	Boolean	No	Lease Scavenging	ipv6_enable_lease_scavenge	When you set this to <b>True</b> , the DHCPv6 server deletes free, expired, and released leases. The default value is <b>False</b> . Example: True

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
v6_leases_scavenging_grace_period	Unsigned Integer	No	Lease Scavenging	ipv6_lease_scavenge_time	Indicates the period (in seconds) for which free, expired, and released DHCPv6 leases remain in the database before they are automatically deleted. Example: 604800

e

## Dynamic Update Group

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-ddnsprincipalgroup	String	Yes			Example: ddnsprincipalgroup
name	String	Yes	Name	name	Example: corp.example.com
comment	String	No	Comment	comment	

## Dynamic Update Cluster Group

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-ddnsprincipalcluster	String	Yes			Example: ddnsprincipalcluster
name	String	Yes	Name	name	Example: cluster1
parent	String	Yes	Dynamic Update Group	clusters	Example: corp.example.com
principals	String	Yes	Principal	principals	Example: WIN-2008-CLUSTER1-NODE1\$@CORP.EXAMPLE.COM
comment	String	No	Comment	comment	

## Grid DHCP Objects

NIOS does not support add and delete operations.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-GridDhcp	String	Yes			Identifies the first row as a header row for the Grid DHCP objects. Example: GridDhcp
authority	Boolean	No			When you set this to <b>True</b> , DHCP server is authoritative for this domain. Example: False
domain_name	String	No			Indicates the domain name.
recycle_leases	Boolean	No	Lease Deletion	recycle_leases	This field is set to <b>True</b> by default. When you set this to <b>True</b> , leases in a deleted range are kept until expiration. Ensure that you use the <b>Overwrite</b> option if you want to change the value to <b>False</b> . Merging data from an import preserves the default value.
ignore_dhcp_option_list_request	Boolean	No		ignore_dhcp_option_list_request	When this is set to <b>True</b> , it clears the value of option-55. Example: False
enable_pxe_lease_time	Boolean	No	Enable PXE lease time		If this value is set to True, DHCP server uses different lease time for PXE clients. This field applies to the host address. Example: False
pxe_lease_time	Unsigned integer	No	PXE Lease Time	pxe_lease_time	Indicates the lease time for PXE clients in seconds. This field applies to the host address. Example: 43220.
bootfile	String	No	Boot File	bootfile	Indicates the boot file name. Example: bootfile1
bootserver	String	No	Boot Server	bootserver	Indicates the boot server. Example: abc.corp100.com

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
nextserver	String	No	Next Server	nextserver	Indicates the next server. Example: blue.domain.com
deny_bootp	Boolean	No	Deny BOOTP Requests	deny_bootp	When this is set to True, it denies BOOTP requests. This field applies to the host address. Example: FALSE
enable_ddns	Boolean	No	Enable DDNS Updates	enable_ddns	Enable or disable dynamic updates via DHCP to DNS server(s). Example: FALSE
ddns_use_option81	Boolean	No	Option 81 Support	ddns_use_option81	Enable or disable option 81 support. Enables <b>always_update_dns field</b> . Example: TRUE
ddns_server_always_updates	Boolean	No		ddns_server_always_updates	When you set this to True, DHCP server will always update DNS.
ddns_generate_host_name	Boolean	No	Generate Hostname	ddns_generate_host_name	When you set this to True, DHCP server will generate a hostname for DNS updates if not sent by client. Example: TRUE
ddns_ttl	Unsigned integer	No	DDNS Update TTL	ddns_ttl	Indicates the DDNS TTL value in seconds. This is an inherited field. Example: 1200
retry_ddns_updates	Boolean	No		retry_ddns_updates	When you set this value to True, DHCP server will retry failed DNS updates. Example: False
ddns_retry_interval	Unsigned integer	No		ddns_retry_interval	Indicates the minimum time in minutes between DNS update retries. You must set <b>ddns_retry_updates</b> to <b>True</b> to modify the <b>ddns_retry_interval</b> value in the CSV file.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
enable_dhcp_thresholds	Boolean	No	Enable DHCP Thresholds	enable_dhcp_thresholds	Enable DHCP thresholds. When you set this field to <b>True</b> , you must enter values in the <b>range_high_water_mark</b> and <b>range_low_water_mark</b> fields. You cannot leave those fields empty. Otherwise, the appliance generates an error.
high_water_mark	Unsigned integer	No	High Water Mark	high_water_mark	Indicates the percentage value for DHCP range usage after which an alarm will be active. When you set <b>enable_thresholds</b> to <b>True</b> , you must enter values in this field and in the <b>range_low_water_mark</b> field. You cannot leave these fields empty. Otherwise, the appliance displays an error message. Example: 80
high_water_mark_reset	Unsigned integer	No			Indicates the percentage value for DHCP range usage after which an alarm will be reset. Example: 85
low_water_mark	Integer	No	Low Water Mark	low_water_mark	Indicates the percentage value for DHCP range usage below which an alarm will be active. When you set <b>enable_thresholds</b> to <b>True</b> , you must enter values in this field and in the <b>range_high_water_mark</b> field. You cannot leave these fields empty. Otherwise, the appliance displays an error message. Example: 10
low_water_mark_reset	Unsigned integer	No			Indicates the percentage value for DHCP range usage above which an alarm will be reset. Example: 10



Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
enable_email_warnings	Boolean	No		enable_email_warnings	Enable to send DHCP threshold warnings via email. Example: False
enable_snmp_warnings	Boolean	No			Enable to send DHCP threshold warnings via SNMP. Example: False
email_list	Email address list	No			List of email addresses. Example: "admin1@infoblox.com", "admin2@somewhere.com",
ipv6_domain_name_servers	IPv6 DNS server list	No			List of IPv6 domain name servers. Example: "2001::1, 2001::2,..."
ping_count	Unsigned integer	No		ping_count	Indicates the number of DHCP pings. Example: 1
ping_timeout	Unsigned integer	No		ping_timeout	Indicates the timeout (in seconds) for DHCP pings. Example: 1000
capture_hostname	Boolean	No			When you set this value to <b>True</b> , the appliance captures host name and lease time when assigning fixed addresses.
enable_leasquery	Boolean	No			When you set this value to <b>True</b> , the appliance allows lease query. Example: False
update_dns_on_lease_renewal	Boolean	No	Lease Renewal Update	override_update_dns_on_lease_renewal	Indicates whether the DHCP server updates DNS when a DHCP lease is renewed. Specify True to enable it or False to disable it.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ipv6_update_dns_on_lease_renewal	Boolean	No			Set or retrieve the <b>ipv6_update_dns_on_lease_renewal</b> flag. This attribute controls whether the DHCP server updates DNS when an IPv6 DHCP lease is renewed. Specify <b>True</b> to enable this feature or <b>False</b> to disable it. The default value is <b>False</b> .
txt_record_handling	String	No			Specifies how DHCP should treat TXT records while performing DNS update. Example: ISC
lease_scavenge_time	Unsigned integer	No			Indicates the Grid level lease_scavenge_time value. If the value is -1, which means this lease scavenge will be disabled. The minimum value would be 7 * 24 * 60 * 60 ( 7 days).
failover_port	Unsigned integer	No			Indicates the failover port number. The default value is 647. The port number must be between 1 and 63999.
enable_fingerprint	Boolean	No			When you set this value to <b>True</b> , fingerprint matching for incoming lease requests will be enabled. Example: False
ipv6_enable_ddns	Boolean	No			Set this value to <b>True</b> to override the value at the Grid level. Set the parameter to <b>False</b> to inherit the settings from the Grid.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ipv6_ddns_enable_option_fqdn	Boolean	No			Indicates whether the FQDN option sent by the client is to be used, or if the server can automatically generate the FQDN. Default value is half of lease time. Example: False
ipv6_ddns_server_always_updates	Boolean	No			Specify True to enable this feature or False to disable it. The default value is False. You must set this to False to update DNS only if requested by the client.
ipv6_generate_hostname	Boolean	No			When you set this value to True, the hostname is generated if it is not sent by the client. Example: False
ipv6_ddns_domainname	String	No			The DDNS domain name in FQDN format. Example: test_domain.com
ipv6_ddns_ttl	Unsigned integer	No			Indicates the member IPv6 DDNS TTL value in seconds. Example: 0
Preferred_lifetime	Integer	No	Preferred Lifetime	preferred_lifetime	Indicates whether the preferred_lifetime value in the DHCP member is used, instead of the Grid default. Example: 6
valid_lifetime	Unsigned integer	No	Valid Lifetime	valid_lifetime	Indicates whether the valid_lifetime value in the Grid Member DHCP is used, instead of the upper-level default. Example: 43200
ipv6_domain_name	String	No			
ipv6_txt_record_handling	String	No			Example: ISC

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ipv6_capture_hostname	Boolean	No			Example: False
ipv6_recycle_leases	Boolean	No			When you set this to True, the leases are kept in recycle bin until one week after expiration. When you set this to False, the leases are irrecoverably deleted. Example: False
ipv6_enable_retry_updates	Boolean	No			When you set this flag to True, the DHCPv6 server retries failed dynamic DNS updates. The default value is True. Example: False
ipv6_retry_updates_interval	Unsigned integer	No			Set the retry interval when the member DHCPv6 server makes repeated attempts to send DDNS updates to a DNS server. The default retry interval is five minutes.
ddns_domainname	String	No	DDNS Domain Name	ddns_domainname	The DDNS domain name in FQDN format. Example: ddns.corp100.com
lease_per_client_settings	String	No			Defines how the server will release the client lease. This field is set to <b>RELEASE_MATCHING_ID</b> by default. Valid values are ONE_LEASE_PER_CLIENT, RELEASE_MATCHING_ID and NEVER_RELEASE.
ignore_client_identifier	Boolean	No			Indicates if the client identifier will be ignored for a DHCP shared network object. When you set this to <b>True</b> , the client identifier will be ignored. Example: False

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
disable_all_nac_filters	Boolean	No			When you set this value to <b>True</b> , NAC filters will be disabled on the Infoblox Grid. Example: False
format_log_option_82	String	No			Select the logging format to either hexadecimal or a decoded string which is human readable.
OPTION-1	String	No	Custom DHCP Options	options	This field applies to the host address. Example: '255.0.0.0' name implies vendor_class='DHCP' (default)
OPTION-XXXX-200	Option information	No	Custom DHCP Options	options	This field applies to the host address. Example: 'dfdfdf' name implies vendor_class='XXXX', optioncode/number 200
v6_leases_scavenging_enabled	Boolean	No	Lease Scavenging	ipv6_enable_leasescavenge	When you set this to <b>True</b> , the DHCPv6 server deletes free, expired, and released leases. The default value is <b>False</b> . Example: True
v6_leases_scavenging_grace_period	Unsigned integer	No	Lease Scavenging	ipv6_lease_scavenging_time	Indicates the period (in seconds) for which free, expired, and released DHCPv6 leases remain in the database before they are automatically deleted. Example: 604800

## Member DNS Objects

NIOS does not support add and delete operations.



### Note

When you export member DNS properties, the CSV file might include the "unbound\_logging\_level" field with "OPERATIONS" as the value. Although this field is only applicable to the IB-4030-10GE appliance and might not apply to your Grid members, you can still perform CSV import using the CSV export file that contains this field without any issues.

Field Name	Data Type Required Associated (Yes/No) GUI Field	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-MemberDns	String	Yes			Identifies the first row as a header row for the member DNS objects. Example: MemberDns
parent	FQDN	Yes			Indicates the parent object. Example: member1.infoblox.com
dns_over_mgmt	Boolean	No			Enable or disable DNS services on the MGMT port. Example: False
dns_over_lan2	Boolean	No			Enable or disable DNS services on the LAN2 port. Example: False
minimal_response	Boolean	No			Enable or disable minimal response of the DNS server. Example: False
forwarders_only	Boolean	No	Use Forwarders Only		Enable use of forwarders only. Example: False
allow_forwarder	IP address list	No			Indicates the list of forwarders.
member_view_nats	integer	No			Indicates the list of views with NAT address used for creating glue records for the view. Example: dns_view1/ INTERFACE/10.10.10.
enable_notify_source_port	integer	No			Enable or disable <b>notify_source_port</b> . Example: False
notify_source_port	Unsigned integer	No			Indicates the notify source port number.
enable_query_source_port	Boolean	No			Enable or disable <b>query_source_port</b> . Example: False
query_source_port	Unsigned integer	No			Indicates the query source port number.
lame_ttl	Unsigned integer	No			Indicates the lame TTL value in seconds. Example: 600

Field Name	Data Type Required Associated (Yes/No) GUI Field	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
auto_sort_views	Boolean	No			Enable or disable DNS views auto-sort. Example: False
member_views	List of Member views	No			Indicates the list of member views. Example: dns_view1, dns_view2,...
allow_transfer	ACL	No	Allow zone transfers to	allow_transfer	List of <b>address_tsig_ac</b> items. Note that you can import the name of a named ACL in this field. Example: NACL1 or "12.0.0.12/Deny,1234::/64/Allow".
excluded_servers	IP address list	No			List of excluded servers for zone transfers.
zone_transfer_format_option	String	No			Indicates the zone transfer format.
recursion_enabled	Boolean No				Indicates the flag to respond to recursive queries. Example: False
allow_query	ACL	No	Allow queries from	allow_query	List of <b>address_tsig_ac</b> items. It can be an IP address, a network entry, Any or a TSIG-/permission. If the first value is not <b>Any</b> or <b>TSIG-</b> , it is assumed to be an IP address or a network entry. Example: 10.0.0.10/Allow, 11.0.0.0/16/Deny, TSIG-foo/xyz/Allow. It can also be a named ACL. Example: NACL1.
allow_recursive_query	ACL	No			List of <b>address_tsig_ac</b> items. It can be an IP address, a network entry, Any or a TSIG-/permission. If the first value is not <b>Any</b> or <b>TSIG-</b> , it is assumed to be an IP address or a network entry. Example: "10.0.0.10/Allow, 11.0.0.0/16/Deny, TSIG-foo/xyz/Allow,..." or it can be a named ACL. Example: "NACL1"
limit_concurrent_recursive_clients	Boolean	No			Enable limit of concurrent recursive client number. Example: False
concurrent_recursive_clients	Unsigned integer	No			Indicates the number of clients allowed to perform concurrent queries. Example: 1000

Field Name	Data Type Required Associated (Yes/No) GUI Field	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
allow_update	ACL	No	Allow updates from	allow_update	List of <b>address_tsig_ac</b> items. It can be an IP address, a network entry, Any or a TSIG-/permission. If the first value is not <b>Any</b> or <b>TSIG-</b> , it is assumed to be an IP address or a network entry. Example: 10.0.0.10/Allow, 11.0.0.0/16/Deny, TSIG-foo/xyz/Allow. It can also be a named ACL. Example: NACL1.
allow_gss_tsig_zone_updates	Boolean	No			Allow GSS-TSIG clients to perform zone updates. Example: False
allow_update_forwarding	Boolean	No	Allow updates from	forward_to	Enable update forwarding for secondary zones. Example: False
enable_custom_root_server	Boolean	No			Indicates the flag to enable custom root servers. Example: False
root_name_servers	Root nameserver list	No			Indicates the list of custom root servers. Example: mm1.test.com/1.1.1.1/,.. The appliance displays an error message if the <b>root_name_servers</b> column has an empty value when the <b>enable_custom_root_server</b> field is set to <b>True</b> in the imported CSV file.
enable_blackhole	Boolean	No			Enable blackhole setting. Example: False
blackhole	ACL	No			Indicates the list of banned addresses. Example: "NACL" or "12.0.0.12/Deny,1234::/64/Allow .."
notify_delay	Unsigned integer	No		notify_delay	This field specifies the seconds of delay the notify messages are sent to the secondaries. The valid value is between 5 and 86400 seconds. Example: 5
enable_nxdomain_redirect	Boolean	No			Enable intercept and redirect nxdomain responses. Example: False
nxdomain_redirect_addresses	IP address list	No			Indicates the list of IPv4 addresses to redirect to for nxdomain responses. Example: "1.1.1.1,2.2.2.2,..."



Field Name	Data Type Required Associated (Yes/No) GUI Field	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
nxdomain_redirect_ttl	Unsigned integer	No			Indicates the NXDOMAIN redirect ttl in seconds. Example: 60
nxdomain_log_query	Boolean	No			If you set this to <b>True</b> , the appliance logs the NXDOMAIN redirections. Example: False
nxdomain_rulesets	Pattern list	No			Indicates the list of ruleset objects that are used for NXDOMAIN redirection. Example: pattern1/MODIFY, pattern2/PASS, ...
enable_blacklist	Boolean	No		enable_blacklist	Enable or disable blacklisting at the Grid level. Example: False
blacklist_redirect_addresses	IP address list	No		blacklist_redirect_addresses	Indicates the list of IPv4 addresses addresses to which the blacklisted queries are redirected. Example: 1.1.1.1,2.2.2.2
blacklist_action	String	No	Action	blacklist_action	Indicates the action to be performed when a domain name matches the pattern defined in an assigned rule. Example: Refuse
blacklist_redirect_ttl	Unsigned integer	No			Indicates the TTL value of synthetic DNS responses resulted by blacklisted queries. Example: 60
blacklist_log_query	Boolean	No		blacklist_log_query	Indicates if blacklisted queries must be logged. Example: False
blacklist_rulesets	List of domain names	No		blacklist_rulesets	Indicates the ruleset objects that are blacklisted at the Grid level. Example: list1.com, list2.com, ...
enable_dns64	Boolean	No			Enable DNS64 synthesis. Example: False
dns64_groups	List of Dns64 groups	No			List of SynthesisGroup objects. Example: dns64_groupA, dns64_groupB, ...

Field Name	Data Type Required Associated (Yes/No) GUI Field	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
max_cached_lifetime	Unsigned integer	No			Indicates the maximum time in seconds a DNS response can be stored in the hardware acceleration cache. You can specify unsigned integer between 60 and 86400. Default value is 86400.
dns_over_v6_mgmt	Boolean	No			Enable or disable DNS services on the IPv6 MGMT port. Example: False
dns_over_v6_lan2	Boolean	No			Enable or disable DNS services on the IPv6 LAN2 port. Example: False
filter_aaaa	String	No			Indicates the type of AAAA filtering for this Grid DNS object. The default value is <b>No</b> . Example: Yes
filter_aaaa_list	ACL	No			Indicates the list of IPv4 addresses and networks from which queries are received. Note that the AAAA filtering is applied to these addresses. Example: "12.0.0.12/Deny,13.0.0.0/8/Allow,..." or "NACL1"
dns_over_v6_lan	Boolean	No			Example: False
copy_xfer_to_notify	Boolean	No			Enable or disable copying of the allowed IP addresses from zone transfer list into also-notify statement in named.conf. Example: False
transfers_in	Unsigned integer	No			Indicates the number of maximum concurrent transfers for the Grid. You can specify unsigned integers between 10 and 100. The default value is 10. Example: 10
transfers_out	Unsigned integer	No			Indicates the number of maximum outbound concurrent zone transfers for the Grid. You can specify unsigned integers between 10 and 100. The default value is 10. Example: 10

Field Name	Data Type Required Associated (Yes/No) GUI Field	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
transfers_per_ns	Unsigned integer	No			Indicates the number of maximum concurrent transfers per member for the Grid. You can specify unsigned integers between two and 100. The default value is two. Example: 2
serial_query_rate	Unsigned integer	No			Indicates the number of maximum concurrent SOA queries per second for the Grid. You can specify unsigned integers between 20 and 100. The default value is 20. Example: 20
max_cache_ttl	Unsigned integer	No			Indicates the maximum time (in seconds) for which the server will cache positive answers. The default value is 604800.
max_ncache_ttl	Unsigned integer	No			Indicates the maximum time (in seconds) for which the server will cache negative (NXDOMAIN) responses. The default value is 10800. The maximum allowed value is 604800.
disable_edns	Boolean	No			Enable or disable EDNS0 support for queries that require recursive resolution. The default value is <b>False</b> .
query_rewrite_enabled	Boolean	No			When this is set to <b>True</b> , query rewrite is enabled at the Grid level. Example: False
ADMGRP-XXXX	String	No	Permissions Admin Group/ Role	permission	ADMGRP-JimSmith is an example of an admin permission of a specific admin group. Example: RW
rpz_drop_ip_rule_enabled	Boolean	No	Ignore RPZ-IP triggers with too small prefix lengths		When this is set to <b>True</b> , DNS server ignores RPZ-IP rules with prefix lengths that are less than the specified prefix length limit. Example: TRUE
rpz_drop_ip_rule_min_prefix_length_ipv4	Unsigned Integer	No	Minimum IPv4 Prefix Length		Indicates the minimum IPv4 prefix length for RPZ-IP triggers. The default value is 29.
rpz_drop_ip_rule_min_prefix_length_ipv6	Unsigned Integer	No	Maximum IPv4 Prefix Length		Indicates the minimum IPv6 prefix length for RPZ-IP triggers. The default value is 112.

Field Name	Data Type Required Associated (Yes/No) GUI Field	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
atc_forwarding_enable	Boolean	No	Enable Recursive Queries Forwarding to BloxOne Threat Defense Cloud		Enable or disable the forwarding of DNS recursive queries to BloxOne Threat Defense Cloud.
atc_forwarding_access_key	String	No	Access Key		API Access Key for the current member.
atc_forwarding_resolver_addresses	IP Address	No	DFP Name Server		IP address of the local DNS resolver.
atc_forwarding_forward_first	Boolean	No	Fallback to a default resolver if ATC does not respond		Option to resolve the DNS query if there is any resolution failure in the BloxOne Threat Defense Cloud.

## Stub Zone

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-StubZone	String	Yes			Identifies the first row as a header row for the stub zones. Example: StubZone
fqdn	FQDN	Yes	Name	zone	This field combines the AAAA record name and the zone name to form the FQDN. Example: <a href="#">aaaa1.corp100.com</a>
view	String	No	DNS View	views	If no view is specified, the default view is used. Example: Default
zone_format	String	Yes	Type		Valid values are <b>FORWARD</b> , <b>IPV4</b> , and <b>IPV6</b> .

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
prefix	String	No	RFC 2317 Prefix	prefix	Prefix is used for reverse-mapping RFC2317 zones only. If you include a prefix in a forward-mapping zone, the appliance ignores the prefix. No error message is generated.
disabled	Boolean	No	Disable	disable	Enable or disable the stub zone. Example: FALSE
comment	String	No	Comment	comment	Example: This is a stub zone.
disable_forwarding	Boolean	No	Do not use forwarders	disable_forwarding	Enable or disable forwarding. Example: False
stub_from	Master Nameserver list	Yes			List of external stub servers. Example: <a href="#">"nm1.test.com/2.2.2.2,..."</a>
stub_members	Member server list	No			List of stub Grid members.
ns_group	String	No		ns_group	Stub member name server group name. Example: stub_ns_group1.
ns_group_external	String	No		external_ns_group	Forward/Stub server name server group name. Example: ext_ns_group1.
EA-Site	String	No	Extensible attribute	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. Example: California.
EA-Users	List	No	Extensible attribute	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. Example: ['Annie', 'John'].

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
ADMGRP-XXXX	String	No	Permissions Admin Group/Role	permission	ADMGRP-JimSmith is an example of an admin permission of a specific admin group. Example: RW

## Authoritative Name Server Group

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-NsGroup	String	Yes			Identifies the first row as a header row for the authoritative name server group objects. Example: AuthoritativeNsGroup
group_name	String	Yes			Indicates the name of the authoritative name server group. Example: ns_group1
_new_group_name	String	No			You can overwrite the group name.
grid primaries	Grid member list and stealth state	No	Grid Primary/ Stealth	primary stealth	List of primary servers of the name server group. The valid format is: "hostname/stealth" Example: "foo.localadmin/False,corp1.com/True,..."
external primaries	Server list	No	External Primary	primary	List of external primary servers. The valid format is: name/ip/stealth/use_2x_tsig/use_tsig/tsig_name/tsig_key. Only name and IP address are required. If <b>stealth</b> is not specified, <b>use_2x_tsig</b> and <b>use_tsig</b> are used and the default value is set to <b>False</b> . Example: "ext1.test.com/1.1.1.1/FALSE,..."

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
external_secondaries	Server list	No	External Secondary	secondaries	List of external secondary servers. The valid format is: name/ip/stealth/use_2x_tsig/use_tsig/tsig_name/tsig_key. Only name and IP address are required. Default values are assumed for <b>stealth</b> , <b>use_2x_tsig</b> and <b>use_tsig</b> . If either <b>use_2x_tsig</b> or <b>use_tsig</b> is <b>True</b> , then <b>tsig_name</b> and <b>tsig_key</b> are required. Example: "sec1.com/1.1.1.1/FALSE/FALSE/FALSE/foo/sdfsf86ew,..."
grid_secondaries	Member server list	No	Grid Secondary	secondaries	List of Grid secondary servers. The valid format is: hostname/stealth/lead/grid_replicate. Only hostname is required. The appliance assumes default value for <b>stealth</b> . Values are not specified for <b>lead</b> and <b>grid_replicate</b> fields. Example: "member1.localdomain/FALSE/TRUE/FALSE,"
is_grid_default	Boolean	No			Set this to <b>True</b> to set this name server group as Grid default, set to <b>False</b> to unset this name server group as Grid default. Example: False
comment	String	No	Comment	comment	Example: This is a authoritative name server group.
EA-Site	String	No	Extensible attribute	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. Example: California.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
EA-Users	String	No	Extensible attribute	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. Example: John.

## Forward-Mapping Zone

Note that to delete a parent zone and the associated subzones, you must add **remove-subzones** column to the CSV export file and set the value to **True**. If you want to delete only the parent zone, then you must set this column value to **False**.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
Header-ForwardZone	String	Yes			Identifies the first row as a header row for the forward zones. Example: ForwardZone
fqdn	FQDN	Yes	Name	zone	This field combines the AAAA record name and the zone name to form the FQDN. Example: aaaa1.corp100.com
view	String	No	DNS View	views	If no view is specified, the default view is used. Example: Default
zone_format	String	Yes	Type		Valid values are <b>FORWARD</b> , <b>IPV4</b> , and <b>IPV6</b> .
prefix	String	No	RFC 2317 Prefix	prefix	Prefix is used for reverse-mapping RFC2317 zones only. If you include a prefix in a forward-mapping zone, the appliance ignores the prefix. No error message is generated.
disabled	Boolean	No	Disable	disable	Enable or disable the forward zone. Example: FALSE
comment	String	No	Comment	comment	Example: This is a Forward zone.



Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
forward_to	Zone forwarder list	Yes	Default Zone Forwarders		List of forwarders for a Forward type zone. Example: <a href="#">fwd1.test.com/1.1.1.1/...</a>
forwarding_servers	Forwarding members list	No	Members		List of forwarding servers. Example: "infoblox.localdomain ...."
forwarders_only	Boolean	No	Use Forwarders Only		Enable use of forwarders only. Example: False
ns_group	String	No		ns_group	Forwarding member name server group name. Example: fwd_ns_group1.
ns_group_external	String	No		external_ns_group	Forward/Stub server name server group name. Example: ext_ns_group1.
EA-Site	String	No	Extensible attribute	extensible_attributes	EA-Site is an example of a predefined extensible attribute. You can add other predefined attributes to the data file. Example: California.
EA-Users	List	No	Extensible attribute	extensible_attributes	EA-Users is an example of a user defined attribute. You can add other user defined attributes to the data file. Example: ['Annie', 'John'].
ADMGRP-XXXX	String	No	Permissions Admin Group/Role	permission	ADMGRP-JimSmith is an example of an admin permission of a specific admin group. Example: RW
disable_ns_generation	Boolean	No	Disable auto-generation of NS records in parent authoritative zone		Determines whether auto-generation of NS records in the parent zone is disabled or not. When this is set to False, the auto-generation is enabled.

## VLAN Objects

The list of VLAN objects are:

- [VLAN View](#)
- [VLAN Range](#)
- [VLAN Object](#)



### Notes

- Scheduled operations (create/update/delete) are not supported for new VLAN objects. However, scheduled assignment of VLAN objects to IPv4/IPv6 Networks is supported.
- CSV export operations for VLAN objects are supported only from Grid Manager and not from WAPI.

## VLAN View

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
name	String	Yes	Name	name	Name of the VLAN view. Example: VLAN_view1
start_vlan_id	Unsigned Integer	Yes	Start VLAN ID	start_vlan_id	Start ID of the VLAN view. The default value is 1.
end_vlan_id	Unsigned Integer	Yes	End VLAN ID	end_vlan_id	End ID of the VLAN view. The default value is 4094.
comment	String	No	Comment	comment	A descriptive comment for the VLAN view.
allow_range_overlapping	Boolean	No	Allow Range Overlapping	allow_range_overlapping	When set to <b>true</b> , the VLAN ranges under the VLAN view can have overlapping IDs. By default, the value is set to <b>false</b> .

## VLAN Range

VLAN view is an import dependency for VLAN range. Ensure that VLAN view data is available in the CSV file while importing data for VLAN range.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
vlan_view	Object Reference	Yes	VLAN View	vlan_view	Reference to the parent VLAN view to which the VLAN range belongs.
name	String	Yes	VLAN Range Name	name	Name of the VLAN range. Example: VLAN_range3
start_vlan_id	Unsigned Integer	Yes	Start VLAN ID	start_vlan_id	Start ID of the VLAN range.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
end_vlan_id	Unsigned Integer	Yes	End VLAN ID	end_vlan_id	End ID of the VLAN range.
comment	String	No	Comment	comment	A descriptive comment for the VLAN range.

## VLAN Object

VLAN view and VLAN range are import dependencies for a VLAN object depending on whether the VLAN object has been added directly to the VLAN view or as part of a VLAN range. Ensure that the relevant VLAN view or VLAN range data is available in the CSV file while importing data for a VLAN object.

Field Name	Data Type	Required (Yes/No)	Associated GUI Field	Associated PAPI Method	Usage and Guidelines
parent	Object Reference	Yes	VLAN Parent	parent	Reference to the parent VLAN view or VLAN range to which the VLAN object belongs.
name	String	Yes	VLAN Name	name	Name of the VLAN object. Example: VLAN30
id	Unsigned Integer	Yes	VLAN ID	id	VLAN ID value.
comment	String	No	Comment	comment	A descriptive comment for the VLAN object.
reserved	Boolean	No	Reserved	reserved	When set to <b>true</b> , the VLAN object can only be assigned manually to the IPAM object. By default, it is set to <b>false</b> .
description	String	No	Description	description	A description for the VLAN object. The description may be used for longer VLAN names.
contact	String	No	Contact	contact	Contact information for the person or team managing or using the VLAN object.
department	String	No	Department	department	Department where the VLAN object is used.

## Importing Multiple Action CSV file

When you import data, you can include multiple actions, such as add, modify, and delete, in one single CSV file. The multiple action CSV import file contains multiple types of objects with its headers and data rows listed in the order of their dependency hierarchy.

The CSV import option supports insert, merge/override and delete operations. To combine these operations together in a single CSV file, you must specify an optional **IMPORT-ACTION** column in the CSV import file. The column value for each data row describes the type of action that the appliance supports for the respective row. The action values include the following: **I** (Insert), **M** (Merge), **O** (Override), **IM** (Insert + Merge), **IO** (Insert + Override), **D** (Delete).

Note that you must specify appropriate values in the **IMPORT-ACTION** column for each row to perform a multiple action CSV import. The appliance performs the respective operation when you specify **I, M, O, D**, in the **IMPORT-ACTION** column. When you specify **IM** or **IO**, the appliance first checks if the corresponding object exists. If it exists, the appliance performs the merge or override operation accordingly. If the object does not exist, you must first perform the insert operation to add the data.

**Note**

When you choose the **Custom** option, the appliance verifies whether the **IMPORT-ACTION** column is present in the imported CSV file. If the **IMPORT-ACTION** column is empty for a data row, then that data row is considered invalid and the appliance displays an error message. If you choose **Continue on Error**, then the CSV import process continues.

In the list of [CSV Supported Objects for Export/Import](#), the objects are listed in the order of their dependency. The least dependent objects are displayed at the top.

### List of all CSV Supported Objects for Export/Import

Object Category	Export/import objects in the order of dependency	Actions that are not supported I - Insert; D - Delete
Global	NamedACL	
	NamedACL item	
	Network view	
	Upgrade group	
	DNS64 Synthesis Group	
	Rulesets	
	Blacklist Rule	
	NXDOMAIN Rule	
	DiscoverySnmpv3Credentials	
	DiscoverySnmpv1v2Credentials	
	DiscoveryCliCredentials	
	SubGridNetwork	D
	FtpUser	
	RirOrganization	
	IpBlockGroup	
	IpBlock	
Member	I, D	

Object Category	Export/Import objects in the order of dependency	Actions that are not supported I - Insert; D - Delete
DNS Objects	GridDns	I, D
	View	
	MemberDns	I, D
	NS Group	
	Authoritative Zone	
	Delegated Zone	
	Bulk Host	
	IPv4 Host Address	
	Host Record	
	IPv6 Host Address	
	A Record	
	AAAA Record	
	CNAME Record	
	DNAME Record	
	MX Record	
	NAPTR Record	
	NS Record	
	PTR Record	
	TXT Record	
	SRV Record	
Forward Zone		
Stub Zone		

Object Category	Export/Import objects in the order of dependency	Actions that are not supported I - Insert; D - Delete
	Response Policy Zone	
	Response Policy Arecord	
	Response Policy AAAArecord	
	Response Policy IP Arecord	
	Response Policy IP AAAArecord	
	Response Policy MXrecord	
	Response Policy NAPTRrecord	
	Response Policy PTRrecord	
	Response Policy SRVrecord	
	Response Policy TXTrecord	
	Response Policy CNAMErecord	
	Response Policy IP Address	
	Response Policy Client IP Address	
	Response Policy IP Address CNAME	
	Response Policy Client IP Address CNAME	
DHCP Objects	GridDhcp	I,D
	MemberDhcp	I,D
	Network Container	
	IPv4 Network	
	IPv6 Network Container	
	IPv6 Network	
	DHCP Failover Associations	

Object Category	Export/Import objects in the order of dependency	Actions that are not supported I - Insert; D - Delete
	IPv4 Shared Network	
	IPv6 Shared Network	
	IPv4 Reserved Range	
	IPv6 Reserved Range	
	IPv4 Fixed Address/Reservation	
	IPv6 Fixed Address	
	IPv4 Option Space	
	IPv4 Option Definition	
	IPv6 Option Space	
	IPv6 Option Definition	
	IPv4 DHCP Range	
	IPv6 DHCP Range	
	DHCP MAC Filter	
	MAC Filter Address Item	
	Option Filter	
	DHCP Fingerprint	
	DHCP Fingerprint Filter	
	Relay Agent Filter	
	NAC Filter	
	Lease	

## Supported Expressions for Search Parameters

Regular expressions are text strings that you use to describe search patterns. You can use the following special characters to define regular expressions for search parameters.

Special character	Purpose	Example	Meaning
()	Defines the scope and precedence of the operator	gr(a e)y	Matches "gray" or "grey".
	Matches either the regular expression before or after the vertical bar	a c	Matches "a" or "c"
.	Matches any single character	.at	Matches any text string ending with "at", such as "hat", "cat", and "bat".
*	Matches the previous regular expression zero or more times	a*bc	Matches zero or multiple occurrences of "a" followed by "bc", such as "bc", "abc", "aabc", "aaabc", and so on.
+	Matches the previous regular expression one or more times	a+bc	Matches one or more occurrences of "a", followed by "bc", such as "abc", "aabc", "aaabc", and so on.
?	Matches the previous regular expression zero or one time	a?bc	Matches zero or one occurrence of "a", followed by "bc", such as "bc" or "abc".
^	Matches the beginning of a text string	^c	Matches any string beginning with "c", such as "cat".
\$	Matches the end of a text string	com\$	Matches any string ending with "com", such as "Infoblox.com".
[ ]	Matches any character specified in the brackets	[03] [abcd] [15a-d]	Matches "0" or "3". Matches "a", "b", "c", or "d". Matches "1", "5", "a", "b", "c", or "d".
[ n-n ]	Matches single characters contained in the specified range, including the start and end points	[0-3] [a-f]	Matches 0, 1, 2, and 3. Matches a, b, c, d, e, and f.
\{m,n\}	Matches the preceding expression at least m but not more than n times.	a\{3,5\}	Matches "aaa", "aaaa", and "aaaaa".



### Note

You can change a special character—such as the period ( . ), asterisk ( \* ), plus sign ( + ), or question mark ( ? ) — into a literal character by prefixing it with a backslash ( \ ). For example, to specify a literal period, asterisk, plus sign, or question mark, use the characters within the following parentheses: ( \. ), ( \\* ), ( \+ ), ( \? ), ( \^ ), ( \\$ ).